

Q200 SERIES™
PROGRAMMER'S MANUAL

QUANTUM CORPORATION

81-45416

REV	A	B	C
DATE	01/87	04/87	11/87
EC	2021	2237	2407

PREFACE

Quantum reserves the right to make changes and/or improvements to its products without incurring any obligation to incorporate such changes or improvements in units previously sold or shipped.

Quantum publishes descriptive data sheets, a Product Description/OEM Manual, and a Technical Reference Manual for each product line. In addition, important changes to a product are conveyed in the form of a Technical Support Bulletin sent to all product customers of record. Those changes which alter manual content are covered by publishing Addenda to the affected manual.

Quantum publications may be requested from your Quantum Sales Representative, or may be ordered from Quantum in the same manner as spare parts.

Publication No. 81-45416, Rev. C

Copyright 1986, 1987, 1988 Quantum Corporation. All rights reserved.
AIRLOCK and DISCACHE are registered trademarks of Quantum Corporation.
Q200 Series is a trademark of Quantum Corporation.
Printed in U.S.A.

TABLE OF CONTENTS

1.0	SCOPE	1
2.0	APPLICABLE DOCUMENTS	1
3.0	GLOSSARY AND CONVENTIONS	1
3.1	GLOSSARY	1
3.2	EDITORIAL CONVENTIONS	2
4.0	INTERFACE CHARACTERISTICS	2
4.1	SCSI BUS	2
4.2	SCSI BUS SIGNALS	5
4.3	SCSI BUS TIMING	7
5.0	LOGICAL CHARACTERISTICS	8
5.1	SCSI BUS PHASES	8
5.1.1	BUS FREE Phase	8
5.1.2	ARBITRATION Phase	8
5.1.3	SELECTION Phase	9
5.1.4	RESELECTION Phase	10
5.1.5	Information Transfer Phases	11
5.1.6	COMMAND Phase	12
5.1.7	DATA Phase	12
5.1.8	STATUS Phase	13
5.1.9	MESSAGE Phase	13
5.1.10	Signal Restrictions Between Phases	14
5.2	SCSI BUS CONDITIONS	14
5.2.1	ATTENTION Condition	14
5.2.2	RESET Condition	14
5.3	SCSI BUS PHASE SEQUENCES	15
5.3.1	Nonarbitrating Systems	15
5.3.2	Arbitrating Systems	15
5.3.3	All Systems	16
5.4	SCSI POINTERS	16
5.5	MESSAGE SYSTEM SPECIFICATION	17
5.5.1	Message Protocol	17
5.5.2	Messages	18
5.5.3	Message Error Recovery Strategy	19
6.0	COMMAND SPECIFICATION	21
6.1	COMMAND DESCRIPTOR BLOCK (CDB)	21
6.1.1	Opcode	21
6.1.2	Logical Unit Number (LUN)	22
6.1.3	Logical Block Address (LBA)	22
6.1.4	Transfer or Allocation Length	22
6.1.5	Vendor Unique (VU)	23
6.1.6	Reserved (R)	23
6.1.7	Flag Bit (F)	23
6.1.8	Link Bit (L)	23
6.1.9	Relative Address (RELADR) Bit	23

TABLE OF CONTENTS(CONT)

	<u>Page</u>
6.2 COMPLETION STATUS BYTE	23
6.3 READING AND WRITING	25
6.3.1 Logical Blocks	25
6.3.2 Logical Block Address Ranges	26
6.3.3 Transferring Data	26
6.4 CONFIGURING THE Q200 Series OPERATING MODE	27
6.4.1 Q200 Series TM Operating Modes	27
6.4.2 Operating Mode Tables	27
6.5 DESCRIPTION OF COMMANDS	28
6.5.1 TEST UNIT READY Command: Opcode = 00 (hex)	28
6.5.2 REZERO UNIT Command: Opcode = 01 (hex)	29
6.5.3 REQUEST SENSE Command: Opcode = 03 (hex)	29
6.5.4 FORMAT UNIT Command: Opcode = 04 (hex)	37
6.5.5 REASSIGN BLOCKS Command: Opcode = 07 (hex)	44
6.5.6 READ Command: Opcode = 08 (hex)	45
6.5.7 WRITE Command: Opcode = 0A (hex)	46
6.5.8 SEEK Command: Opcode = 0B (hex)	47
6.5.9 INQUIRY Command: Opcode = 12 (hex)	48
6.5.10 MODE SELECT Command: Opcode = 15 (hex)	49
6.5.11 RESERVE Command: Opcode = 16 (hex)	60
6.5.12 RELEASE Command: Opcode = 17 (hex)	61
6.5.13 COPY Command: Opcode = 18 (hex)	62
6.5.14 MODE SENSE Command: Opcode = 1A (hex)	67
6.5.15 START/STOP UNIT Command: Opcode = 1B (hex)	73
6.5.16 SEND DIAGNOSTIC Command: Opcode = 1D (hex)	74
6.5.17 READ CAPACITY Command: Opcode = 25 (hex)	75
6.5.18 READ EXTENDED Command: Opcode = 28 (hex)	76
6.5.19 WRITE EXTENDED Command: Opcode = 2A (hex)	77
6.5.20 SEEK EXTENDED Command: Opcode = 2B (hex)	77
6.5.21 VERIFY Command: Opcode = 2F (hex)	78
6.5.22 READ DEFECT DATA Command: Opcode = 37 (hex)	78
6.5.23 WRITE DATA BUFFER Command: Opcode = 3B (hex)	80
6.5.24 READ DATA BUFFER Command: Opcode = 3C (hex)	81

APPENDIX

A Media Defect Handling	A-1
B DisCache	B-1

LIST OF FIGURES

Figure 4-1 SCSI ID Bits	3
Figure 4-2 Sample SCSI Configuration	4
Figure 5-1 Phase Sequences without Arbitration	15
Figure 5-2 Phase Sequences with Arbitration	16
Figure 5-3 Simplified SCSI System	16
Figure A-1 Factory Found Defects	A-2
Figure A-2 Field Found Defects	A-3

LIST OF TABLES

	<u>Page</u>
Table 4-1 Signal Sources	7
Table 5-1 Information Transfer Phases	12
Table 5-2 Message Codes Supported by the Q200 Series	17
Table 6-1 Commands Supported by the Q200 Series	22
Table 6-2 Status Byte	23
Table 6-3 Status Byte Code Bit Values	24
Table 6-4 Q200 Series Status Codes	25
Table 6-5 TEST UNIT READY Command	28
Table 6-6 REZERO UNIT Command	28
Table 6-7 REQUEST SENSE Command	28
Table 6-8 Extended Sense Format	30
Table 6-9 Sense Error Codes	32
Table 6-10 FORMAT UNIT Command	37
Table 6-11 FORMAT UNIT Command Variations	38
Table 6-12 Defect List - Logical Block Format	39
Table 6-13 REASSIGN BLOCKS Command	44
Table 6-14 Reassign Blocks Defect List	45
Table 6-15 READ Command	45
Table 6-16 WRITE Command	46
Table 6-17 SEEK Command	47
Table 6-18 INQUIRY Command	48
Table 6-19 Inquiry Data Format	48
Table 6-20 MODE SELECT Command	50
Table 6-21 Mode Select Parameter List	51
Table 6-22 RESERVE Command	60
Table 6-23 RELEASE Command	61
Table 6-24 COPY Command	62
Table 6-25 Copy Parameter List	62
Table 6-26 Copy Functions	63
Table 6-27 Segment Descriptor for Copy Function Codes 00H, 01H	64
Table 6-28 Segment Descriptor for Copy Function Code 02H	64
Table 6-29 MODE SENSE Command	67
Table 6-30 Mode Sense Data	69
Table 6-31 START/STOP UNIT Command	73
Table 6-32 SEND DIAGNOSTIC Command	74
Table 6-33 READ CAPACITY Command	75
Table 6-34 Read Capacity Data	75
Table 6-35 READ EXTENDED Command	76
Table 6-36 WRITE EXTENDED Command	77
Table 6-37 SEEK EXTENDED Command	77

LIST OF TABLES (CONT)

	<u>Page</u>
Table 6-38 VERIFY Command	78
Table 6-39 READ DEFECT DATA Command	78
Table 6-40 Defect List Header	79
Table 6-41 Physical Sector Format	80
Table 6-42 Bytes Offset From Index Format	80
Table 6-43 WRITE BUFFER Command	80
Table 6-44 READ BUFFER Command	81
Table 6-45 Read Buffer Header	81
Table B-1 Programmable Parameters	B-2

1.0 SCOPE

This document defines the functional requirements for interfacing with the Q200 Series Intelligent Disk Drives via the Small Computer System Interface (SCSI). In addition to specifying the command set, this document is intended to provide all other information on the Q200 logical interface (e.g., bus timing, message protocol, etc.) necessary for software designers to quickly integrate Q200 Series disk drives into an SCSI host system.

Mechanical and electrical parameters are not addressed by this document. Refer to the ANSI X3.131-1986 Small Computer System Interface specification for information on these parameters.

2.0 APPLICABLE DOCUMENTS

ANSI X3.131-1986 Small Computer System Interface specification

Common Command Set Rev 4B (by an ANSC X3T9.2 working group)

Q250 / Q280 Disk Drive OEM Manual (P/N 81-45233)

3.0 GLOSSARY AND CONVENTIONS

3.1 GLOSSARY

byte - This term indicates an 8-bit (octet) byte.

command descriptor block (CDB) - The structure used to communicate requests from an initiator to a target.

connect - The function that occurs when an initiator selects a target to start an operation.

disconnect - The function that occurs when a target releases control of the SCSI bus, allowing it to go to the BUS FREE phase.

initiator - An SCSI device (usually a host system) that requests an operation to be performed by another SCSI device.

logical unit - A physical or virtual device addressable through a target.

logical unit number - An encoded three-bit identifier for the logical unit.

LSB - Least Significant Bit

LUN - Logical Unit Number

ms - Millisecond

MSB - Most Significant Bit

ns - Nanosecond

one - A true signal value.

peripheral device - A peripheral that can be attached to an SCSI device (e.g., magnetic-disk, printer, optical-disk, or magnetic-tape).

reconnect - The function that occurs when a target selects an initiator to continue an operation after a disconnect.

reserved - The term used for bits, bytes, fields, and code values that are set aside for future standardization.

SCSI address - The octal representation of the unique address (0-7) assigned to an SCSI device. This address would normally be assigned and set in the SCSI device during system installation.

SCSI ID - The bit-significant representation of the SCSI address referring to one of the signal lines DB(7-0).

SCSI device - A host computer adapter, a peripheral controller or an intelligent peripheral that can be attached to the SCSI bus.

signal assertion - The act of driving a signal to the true state.

signal negation - The act of driving a signal to the false state or allowing the cable terminators to bias the signal to the false state (by placing the driver in the high impedance condition).

signal release - The act of allowing the cable terminators to bias the signal to the false state (by placing the driver in the high impedance condition).

status - One byte of information sent from a target to an initiator upon completion of each command.

target - An SCSI device that performs an operation requested by an initiator.

us - Microsecond

vendor unique - This term indicates bits, fields, or code values that are vendor specific.

xx_B - Numbers followed by a capital B subscript are binary values.

xx_H - Numbers followed by a capital H subscript or "(hex)" are hexadecimal values. All other numbers are decimal values.

zero - A false signal value.

3.2 EDITORIAL CONVENTIONS

Certain words and terms used in this manual have a specific meaning beyond the normal English meaning. These words and terms are defined either in the glossary (see 3.1) or in the text where they first appear (e.g., arbitration delay - see 4.3.1). Names of signals, phases, conditions, messages, commands, statuses, and sense keys are all in uppercase (e.g., REQUEST SENSE). Lowercase is used for words having the normal English meaning.

Throughout this manual, the term Q200 will be used to refer to either the Q250 or the Q280 disk drive.

4.0 INTERFACE CHARACTERISTICS

4.1 SCSI BUS

Communication on the SCSI bus is allowed between only two SCSI devices at any given time. There is a maximum of eight SCSI devices. Each SCSI device has an SCSI ID bit assigned as shown in Figure 4-1.

When two SCSI devices communicate on the SCSI bus, one acts as an initiator and the other acts as a target. The initiator originates an operation and the target performs the operation. An SCSI device usually has a fixed role as

either an initiator or a target, but some devices may be able to assume either role. The Q200 acts only as an initiator when executing COPY commands with a third party device (e.g., backing up Q200 data to a tape file).

An initiator may address up to eight peripheral devices that are connected to a target. Three sample system configurations are shown in Figure 4-2.

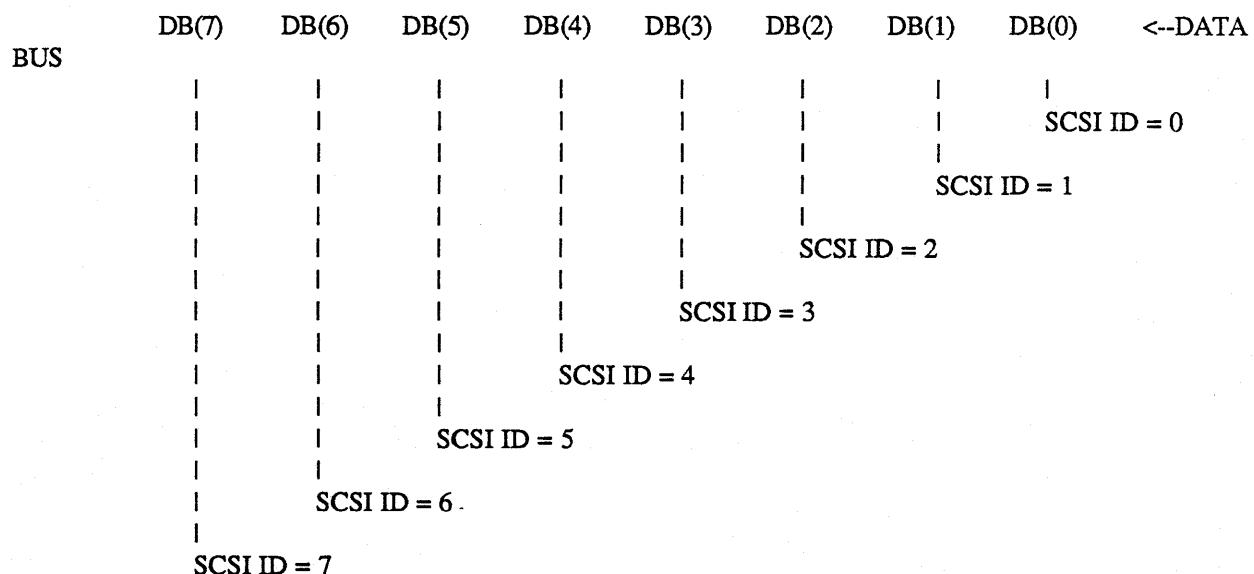
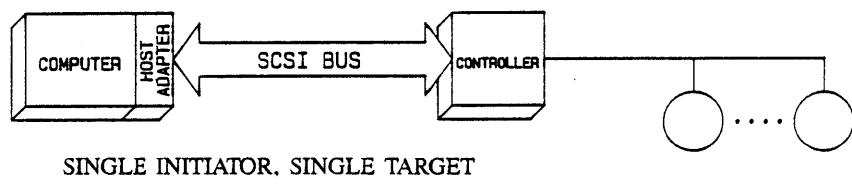
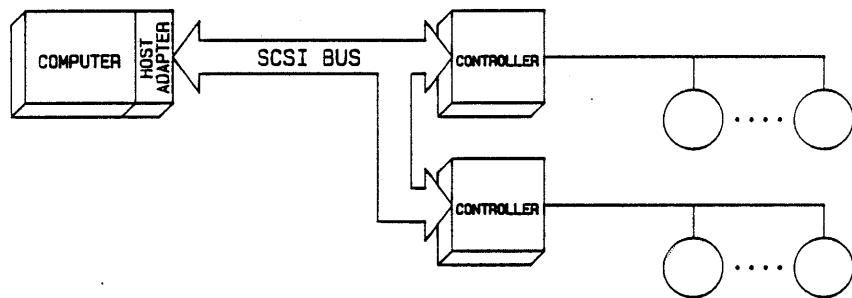


FIGURE 4-1. SCSI ID Bits

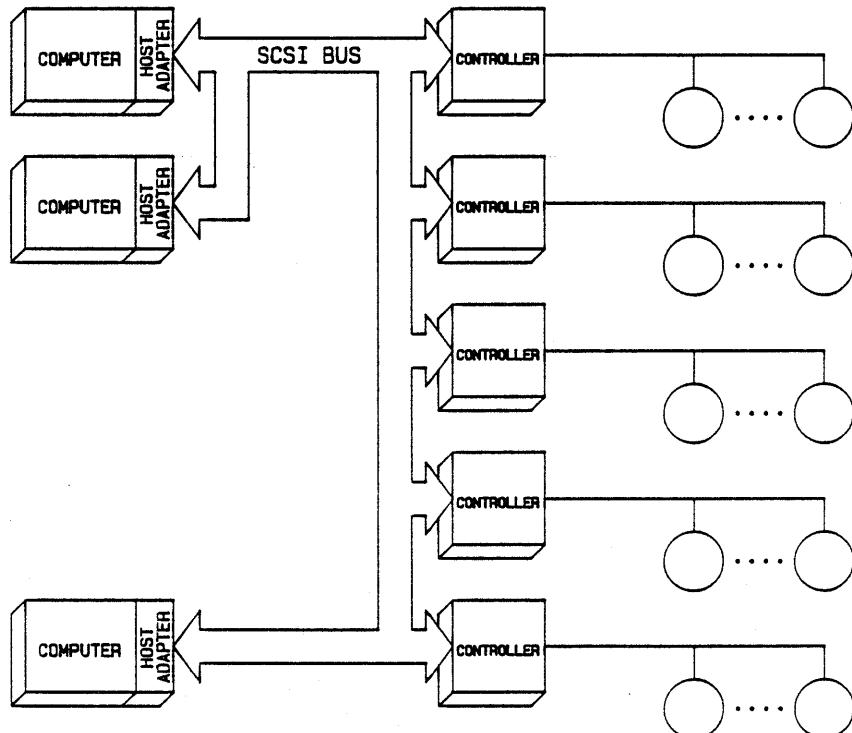
Peripheral devices such as magnetic disks, printers, optical-disks, and magnetic-tapes.



SINGLE INITIATOR, SINGLE TARGET



SINGLE INITIATOR, MULTI TARGET



MULTI INITIATOR, MULTI TARGET

FIGURE 4-2. Sample SCSI Configuration

Up to eight SCSI devices can be supported on the SCSI bus. They can be any combination of initiators and targets.

Certain SCSI bus functions are assigned to either the initiator or to the target. The initiator may arbitrate for the SCSI bus and select a particular target. The target may request the transfer of COMMAND, DATA, STATUS, or other information on the DATA BUS, and in some cases it may arbitrate for the SCSI bus and reselect an initiator for the purpose of continuing an operation.

Information transfers on the DATA BUS are asynchronous and follow a defined REQ/ACK handshake protocol. One byte of information may be transferred with each handshake. The synchronous data transfer option is not supported on the Q200.

4.2 SCSI BUS SIGNALS

There are a total of eighteen signals. Nine are used for control and nine are used for data. (Data signals include the parity signal option.) These signals are described as follows:

BSY (BUSY) - An "OR-tied" signal that indicates that the bus is being used.

SEL (SELECT) - A signal used by an initiator to select a target or by a target to reselect an initiator.

C/D (CONTROL/DATA) - A signal driven by a target that indicates whether CONTROL or DATA information is on the DATA BUS. True indicates CONTROL.

I/O (INPUT/OUTPUT) - A signal driven by a target that controls the direction of data movement on the DATA BUS with respect to an initiator. True indicates input to the initiator. This signal is also used to distinguish between SELECTION and RESELECTION phases.

MSG (MESSAGE) - A signal driven by a target during the MESSAGE phase.

REQ (REQUEST) - A signal driven by a target to indicate a request for a REQ/ACK data transfer handshake.

ACK (ACKNOWLEDGE) - A signal driven by an initiator to indicate an acknowledgment for a REQ/ACK data transfer handshake.

ATN (ATTENTION) - A signal driven by an initiator to indicate the ATTENTION condition.

RST (RESET) - An "OR-tied" signal that indicates the RESET condition.

DB(7-0,P) (DATA BUS) - Eight data-bit signals, plus a parity-bit signal that form a DATA BUS. DB(7) is the most significant bit and has the highest priority during the ARBITRATION phase. Bit number, significance, and priority decrease downward to DB(0). A data bit is defined as one when the signal value is true and is defined as zero when the signal value is false.

DB(P) is set to either zero or one so that data parity is odd. The use of parity is a system option, but parity generation and checking must be consistent for all SCSI devices on the bus (i.e., either all SCSI devices on the bus must generate and check parity, or all SCSI devices on the bus must not check parity). Parity is not valid during the ARBITRATION phase.

4.2.1 SIGNAL VALUES

Signals may assume true or false values. If any driver is asserted, then the signal is true. Whenever the term negated is used in this manual, it means that the signal may, at the option of the implementor, be actively driven false or may be simply released (in which case the bias circuitry of the bus terminators pulls the signal false

whenever it is released by the drivers at every SCSI device). For the Q200, negated signals are released and asserted signals are actively driven true.

4.2.2 OR-TIED SIGNALS

The BSY and RST signals shall be OR-tied only. In the ordinary operation of the bus, these signals are simultaneously driven true by several drivers. No signals other than BSY, RST, and DB(P) are simultaneously driven by two or more drivers, and any signal other than BSY and RST may employ OR-tied or non-OR-tied drivers. DB(P) shall not be driven false during the ARBITRATION phase. There is no operational problem in mixing OR-tied and non-OR-tied drivers on signals other than BSY and RST.

4.2.3 SIGNAL SOURCES

Table 4-1 indicates which type of SCSI device is allowed to source each signal. No attempt is made to show if the source is driving asserted, driving negated, or is passive. All SCSI device drivers that are not active sources shall be in the passive state. Note that the RST signal may be sourced by any SCSI device at any time.

Table 4-1
SIGNAL SOURCES

Signals					
Bus Phase	BSY	SEL	C/D, I/O MSG, REQ	ACK/ATN	DB(7-0,P)
BUS FREE	None	None	None	None	None
ARBITRATION	All	Winner	None	None	SCSI ID
SELECTION	I&T	Initiator	None	Initiator	Initiator
RESELECTION	I&T	Target	Target	Initiator	Target
COMMAND	Target	None	Target	Initiator	Initiator
DATA IN	Target	None	Target	Initiator	Target
DATA OUT	Target	None	Target	Initiator	Initiator
STATUS	Target	None	Target	Initiator	Target
MESSAGE IN	Target	None	Target	Initiator	Target
MESSAGE OUT	Target	None	Target	Initiator	Initiator

All: The signal shall be driven by all SCSI devices that are actively arbitrating.

SCSI ID: A unique data bit (the SCSI ID) shall be driven by each SCSI device that is actively arbitrating; the other seven data bits shall be released (i.e., not driven) by this SCSI device. The parity bit (DP(P)) may be undriven or driven to the true state, but shall never be driven to the false state during this phase.

I&T: The signal shall be driven by the initiator, target, or both, as specified in the SELECTION phase and RESELECTION phase.

Initiator: If this signal is driven, it shall be driven only by the active initiator.

None: The signal shall be released; that is, not driven by any SCSI device. The bias circuitry of the bus terminators pulls the signal to the false state.

Winner: The signal shall be driven by the one SCSI device that wins arbitration.

Target: If the signal is driven, it shall be driven only by the active target.

4.3 SCSI BUS TIMING

Unless otherwise indicated, the delay-time measurements for each SCSI device, shown in Sections 4.3.1 through 4.3.11, shall be calculated from signal conditions existing at that SCSI device's own SCSI bus connection. Thus, these measurements (except cable skew delay) can be made without considering delays in the cable.

4.3.1 ARBITRATION DELAY (2.2 microseconds)

The minimum time an SCSI device shall wait from asserting BSY for arbitration until the DATA BUS can be examined to see if arbitration has been won. There is no maximum time.

4.3.2 BUS CLEAR DELAY (800 nanoseconds)

The maximum time for an SCSI device to stop driving all bus signals after:

- 1) The BUS FREE phase is detected (BSY and SEL both false for a bus settle delay).
- 2) SEL is received from another SCSI device during the ARBITRATION phase.
- 3) The transition of RST to true.

NOTE: For the first condition above, the minimum time for an SCSI device to clear the bus is 1200 nanoseconds from BSY and SEL first becoming both false. If an SCSI device requires more than a bus settle delay to detect BUS FREE phase, it shall clear the bus within a bus clear delay minus the excess time.

4.3.3 BUS FREE DELAY (800 nanoseconds)

The minimum time that an SCSI device shall wait from its detection of the BUS FREE phase (both BSY and SEL are false for a bus settle delay) until its assertion of BSY when going into the ARBITRATION phase.

4.3.4 BUS SET DELAY (1.8 microseconds)

The maximum time before an SCSI device must assert BSY and its SCSI ID bit on the DATA BUS after it detects BUS FREE phase (both BSY and SEL are false for a bus settle delay) for the purpose of entering the ARBITRATION phase.

4.3.5 BUS SETTLE DELAY (400 nanoseconds)

The time to wait for the bus to settle after changing certain control signals as called out in the protocol definitions.

4.3.6 CABLE SKEW DELAY (10 nanoseconds)

The maximum difference in propagation time allowed between any two SCSI bus signals when measured between any two SCSI devices.

4.3.7 DATA RELEASE DELAY (400 nanoseconds)

The maximum time for an initiator to release the DATA BUS signals following the transition of the I/O signal from false to true.

4.3.8 DESKEW DELAY (45 nanoseconds)

The minimum time required for deskew of certain signals.

4.3.9 RESET HOLD TIME (25 microseconds)

The minimum time for which RST is asserted. There is no maximum time.

4.3.10 SELECTION ABORT TIME (200 microseconds)

The maximum time that a target (or initiator) shall take from its most recent detection of being selected (or reselected) until asserting a BSY response. This timeout is required to ensure that a target (or initiator) does not assert BSY after a SELECTION (or RESELECTION) phase has been aborted. This is not the selection timeout period; see Sections 5.1.3.5 and 5.1.4.2 for a complete description.

4.3.11 SELECTION TIMEOUT DELAY (250 milliseconds)

The minimum time that an initiator (or target) should wait for a BSY response during the SELECTION (or RESELECTION) phase before starting the timeout procedure.

5.0 LOGICAL CHARACTERISTICS

5.1 SCSI BUS PHASES

The SCSI architecture includes eight distinct phases:

BUS FREE phase
ARBITRATION phase
SELECTION phase
RESELECTION phase
COMMAND phase
DATA phase
STATUS phase
MESSAGE phase

These phases are collectively termed the
information transfer phases.

The SCSI bus can never be in more than one phase at any given time. Unless otherwise noted in the following descriptions, signals that are not mentioned shall not be asserted.

5.1.1 BUS FREE PHASE

The BUS FREE phase is used to indicate that no SCSI device is actively using the SCSI bus and that the SCSI bus is available for subsequent users.

SCSI devices shall detect the BUS FREE phase after SEL and BSY are both false for at least a bus settle delay.

SCSI devices shall release all SCSI bus signals within a bus clear delay after BSY and SEL become continuously false for a bus settle delay. If an SCSI device requires more than a bus settle delay to detect the BUS FREE phase, then it shall release all SCSI bus signals within a bus clear delay minus the excess time to detect the BUS FREE phase. The total time to clear the SCSI bus shall not exceed a bus settle delay plus a bus clear delay.

5.1.2 ARBITRATION PHASE

The ARBITRATION phase allows one SCSI device to gain control of the SCSI bus so that it can assume the role of an initiator or target.

NOTE: Implementation of the ARBITRATION phase is a system option. Systems that do not implement this option can have only one initiator. The ARBITRATION phase is required for systems that use the RESELECTION phase.

The procedure for an SCSI device to obtain control of the SCSI bus is as follows:

- 1) The SCSI device shall first wait for the BUS FREE phase to occur. The BUS FREE phase is detected whenever both BSY and SEL are simultaneously and continuously false for a minimum of a bus settle delay. (Implementor's NOTE: This bus settle delay is necessary because a transmission line phenomenon known as a "wire-OR glitch" may cause BSY to briefly appear false, even though it is being driven true.)
- 2) The SCSI device shall wait a minimum of a bus free delay after detection of the BUS FREE phase (i.e., after BSY and SEL are both false for a bus settle delay) before driving any signal.
- 3) Following the bus free delay in Step 2, the SCSI device may arbitrate for the SCSI bus by asserting both BSY and its own SCSI ID; however, the SCSI device shall not start to arbitrate (i.e., assert BSY and its SCSI ID) if more than a bus set delay has passed since the BUS FREE phase was last observed. (Implementor's NOTE: There is no maximum delay before asserting BSY and the SCSI ID following the bus free delay in Step 2 as long as the bus remains in the BUS FREE phase. However, SCSI devices that delay longer than a bus settle delay plus a bus set delay from the time when BSY and SEL first become false may fail to participate in arbitration when competing with faster SCSI devices.)
- 4) After waiting at least an arbitration delay (measured from its assertion of BSY) the SCSI device shall examine the DATA BUS. If a higher priority SCSI ID bit is true on the DATA BUS (DB(7) is the highest), then the SCSI device has lost the arbitration and the SCSI device must release its signals and return to Step 1. If no higher priority SCSI ID bit is true on the DATA BUS, then the SCSI device has won the arbitration and it shall assert SEL. Any other SCSI device that is participating in the ARBITRATION phase has lost the arbitration and shall release BSY and its SCSI ID bit within a bus clear delay after SEL becomes true. An SCSI device that loses arbitration may return to Step 1.
- 5) The SCSI device that wins arbitration shall wait at least a bus clear delay plus a bus settle delay after asserting SEL before changing any signals.

NOTE: The SCSI ID bit is a single bit on the DATA BUS that corresponds to the SCSI device's unique SCSI address. All other seven DATA BUS bits shall be released by the SCSI device. Parity is not valid during the ARBITRATION phase. During the ARBITRATION phase, DB(P) may be undriven or driven to the true state, but shall not be driven to the false state.

5.1.3 SELECTION PHASE

The SELECTION phase allows an initiator to select a target for the purpose of initiating some target function (e.g., READ or WRITE command).

NOTE: The I/O signal shall be negated so that this phase can be distinguished from RESELECTION phase.

5.1.3.1 NONARBITRATING SYSTEMS

In systems with the ARBITRATION phase not implemented, the initiator shall first detect the BUS FREE phase and then wait a minimum of a bus clear delay. Then, except in certain single initiator environments with initiators employing the single initiator option (see 5.1.3.4), the initiator shall assert the desired target's SCSI ID and its own initiator SCSI ID (and the parity bit in systems with parity implemented) on the DATA BUS. After two deskelw delays, the initiator shall assert SEL.

5.1.3.2 ARBITRATING SYSTEMS

In systems with ARBITRATION phase implemented, the SCSI device that won the arbitration has both BSY and SEL asserted and has delayed at least a bus clear delay plus a bus settle delay before ending the ARBITRATION phase. The SCSI device that won the arbitration becomes an initiator by not asserting I/O. Except in certain single initiator environments with initiators employing the single initiator option (see 5.1.3.4), the initiator shall assert the desired target's SCSI ID and its own initiator SCSI ID (and the parity bit in systems with parity implemented) on the DATA BUS. The initiator shall then wait at least two deskew delays and release BSY. The initiator shall then wait at least a bus settle delay before looking for a response from the target.

5.1.3.3 ALL SYSTEMS

In all systems, the target shall determine that it is selected when SEL and its SCSI ID bit are true and BSY and I/O are false for at least a bus settle delay. The selected target may examine the DATA BUS in order to determine the SCSI ID of the selecting initiator unless the initiator employed the single initiator option (see 5.1.3.4). The selected target shall then assert BSY within a selection abort time of its most recent detection of being selected; this is required for correct operation of the timeout procedure. In systems with parity implemented, the target shall not respond to a selection if bad parity is detected. Also, if more than two SCSI ID bits are on the DATA BUS, the target shall not respond to selection.

At least two deskew delays after the initiator detects BSY is true, it shall release SEL and may change the DATA BUS.

5.1.3.4 SINGLE INITIATOR OPTION

Initiators that do not implement the RESELECTION phase and do not operate in the multiple initiator environment are allowed to set only the target's SCSI ID bit and not the initiator's SCSI ID bit during the SELECTION phase. This makes it impossible for the target to determine the initiator's SCSI ID.

5.1.3.5 SELECTION TIMEOUT PROCEDURE

Two optional selection timeout procedures are specified for clearing the SCSI bus if the initiator waits a minimum of a selection timeout delay and there has been no BSY response from the target:

- 1) Optionally, the initiator shall assert the RST signal (see 5.2.2).
- 2) Optionally, the initiator shall continue asserting SEL and shall release the DATA BUS. If the initiator has not detected BSY to be true after at least a selection abort time plus two deskew delays, the initiator shall release SEL allowing the SCSI bus to go to the BUS FREE phase. SCSI devices shall ensure that when responding to selection that the selection was still valid within a selection abort time of their assertion of BSY. Failure to comply with this requirement could result in an improper selection (two targets connected to the same initiator, wrong target connected to an initiator, or a target connected to no initiator).

5.1.4 RESELECTION PHASE (OPTIONAL)

RESELECTION is an optional phase that allows a target to reconnect to an initiator for the purpose of continuing some operation that was previously started by the initiator but was suspended by the target (i.e., the target disconnected by allowing a BUS FREE phase to occur before the operation was complete).

5.1.4.1 RESELECTION

RESELECTION can only be used in systems that have ARBITRATION phase implemented.

Upon completing the ARBITRATION phase, the winning SCSI device has both BSY and SEL asserted and has delayed at least a bus clear delay plus a bus settle delay. The winning SCSI device becomes a target by asserting the I/O signal. The winning SCSI device shall also set the DATA BUS to a value that is the OR of its SCSI ID bit and the initiator's SCSI ID bit. The target shall wait at least two deskew delays and release BSY. The target shall then wait at least a bus settle delay before looking for a response from the initiator.

The initiator shall determine that it is reselected when SEL, I/O, and its SCSI ID bit are true and BSY is false for at least a bus settle delay. The reselected initiator may examine the DATA BUS in order to determine the SCSI ID of the reselecting target. The reselected initiator shall then assert BSY within a selection abort time of its most recent detection of being reselected; this is required for correct operation of the timeout procedure. In systems with parity implemented, the initiator shall not respond to a RESELECTION if bad parity is detected. Also, the initiator shall not respond to a RESELECTION if more than two SCSI ID bits are on the DATA BUS.

After the target detects BSY, it shall also assert BSY and wait at least two deskew delays and then release SEL. The target may then change the I/O signal and the DATA BUS. After the reselected initiator detects SEL false, it shall release BSY. The target shall continue asserting BSY until the target is ready to relinquish the SCSI bus.

NOTE: When the target is asserting BSY, a transmission line phenomenon known as a "wire-OR-glitch" may cause BSY to appear false for up to a round-trip propagation delay following the release of BSY by the initiator. This is the reason why the BUS FREE phase is recognized only after both BSY and SEL are continuously false for a minimum of a bus settle delay. Cables longer than 25 meters should not be used even if the chosen driver, receiver and cable provide adequate noise margins because they increase the duration of the glitch and could cause SCSI devices to inadvertently detect the BUS FREE phase.

5.1.4.2 RESELECTION TIMEOUT PROCEDURE

Two optional RESELECTION timeout procedures are specified for clearing the SCSI bus during a RESELECTION phase if the target waits for a minimum of a selection timeout period and there has been no BSY response from the initiator:

- 1) Optionally, the target shall assert the RST signal (see 5.2.2).

NOTE: The Q200 will never assert the RST signal.

- 2) Optionally, the target shall continue asserting SEL and I/O and shall release all DATA BUS signals. If the target has not detected BSY to be true after at least a selection abort time plus two deskew delays, the target shall release SEL and I/O allowing the SCSI bus to go to the BUS FREE phase. SCSI devices that respond to RESELECTION shall ensure that the RESELECTION was still valid within a selection abort time of their assertion of BSY. Failure to comply with this requirement could result in an improper reselection (two initiators connected to the same target or the wrong initiator connected to a target).

5.1.5 INFORMATION TRANSFER PHASES

NOTE: The COMMAND, DATA, STATUS, and MESSAGE phases are all grouped together as the information transfer phases because they are all used to transfer data or control information via the DATA BUS. The actual contents of the information is beyond the scope of this section.

The C/D, I/O, and MSG signals are used to distinguish between the different information transfer phases. (See Table 5-1.) The target drives these three signals and, therefore, controls all changes from one phase to another. The initiator can request a MESSAGE OUT phase by asserting ATN, while the target can cause the BUS FREE phase by releasing MSG, C/D, I/O, and BSY.

Table 5-1
INFORMATION TRANSFER PHASES

Signal					
MSG	C/D	I/O	Phase Name	Direction of Transfer	Comment
0	0	0	DATA OUT	Initiator to target	Data
0	0	1	DATA IN	Initiator from target	Phase
0	1	0	COMMAND	Initiator to target	
0	1	1	STATUS	Initiator from target	
1	0	0	*		
1	0	1	*		
1	1	0	MESSAGE OUT	Initiator to target	Message
1	1	1	MESSAGE IN	Initiator from target	Message

Key: 0 = False, 1 = True, * = Reserved for future standardization

The information transfer phases use one or more REQ/ACK handshakes to control the information transfer. Each REQ/ACK handshake allows the transfer of one byte of information. During the information transfer phases BSY shall remain true and SEL shall remain false. Additionally, during the information transfer phases, the target shall continuously envelope the REQ/ACK handshake(s) with C/D, I/O, and MSG in such a manner that these control signals are valid for a bus settle delay before the assertion of REQ of the first handshake and remain valid until the negation of ACK at the end of the last handshake.

The target shall control the direction of information transfer by means of the I/O signal. When I/O is true, information shall be transferred from the target to the initiator. When I/O is false, information shall be transferred from the initiator to the target.

If I/O is true (transfer to the initiator), the target shall first drive DB(7-0,P) to their desired values, delay at least one deskew delay plus a cable skew delay, then assert REQ. DB(7-0,P) shall remain valid until ACK is true at the target. The initiator shall read DB(7-0,P) after REQ is true, then signal its acceptance of the data by asserting ACK. When ACK becomes true at the target, the target may change or release DB(7-0,P) and shall negate REQ. After REQ is false the initiator shall then negate ACK. After ACK is false, the target may continue the transfer by driving DB(7-0,P) and asserting REQ, as described above.

If I/O is false (transfer to the target) the target shall request information by asserting REQ. The initiator shall drive DB(7-0,P) to their desired values, delay at least one deskew delay plus a cable skew delay and assert ACK. When ACK becomes true at the target, the target shall read DB(7-0,P), then negate REQ. When REQ becomes false at the initiator, the initiator may change or release DB(7-0,P) and shall negate ACK. The target may continue the transfer by asserting REQ, as described above.

5.1.6 COMMAND PHASE

The COMMAND phase allows the target to request command information from the initiator.

The target shall assert the C/D signal and negate the I/O and MSG signals during the REQ/ACK handshake(s) of this phase.

5.1.7 DATA PHASE

The data phase is a term that encompasses both the DATA IN phase and the DATA OUT phase.

5.1.7.1 DATA IN PHASE

The DATA IN phase allows the target to request that data be sent to the initiator from the target.

The target shall assert the I/O signal and negate the C/D and MSG signals during the REQ/ACK handshake(s) of this phase.

5.1.7.2 DATA OUT PHASE

The DATA OUT phase allows the target to request that data be sent from the initiator to the target.

The target shall negate the C/D, I/O, and MSG signals during the REQ/ACK handshake(s) of this phase.

5.1.8 STATUS PHASE

The STATUS phase allows the target to request that status information be sent from the target to the initiator.

The target shall assert C/D and I/O and negate the MSG signal during the REQ/ACK handshake of this phase.

5.1.9 MESSAGE PHASE

The message phase is a term that references either a MESSAGE IN, or a MESSAGE OUT phase. The first byte transferred in either of these phases shall be either a single-byte message or the first byte of a multiple-byte message. Multiple-byte messages shall be wholly contained within a single message phase.

NOTE: The Q200 does not support multiple-byte messages.

5.1.9.1 MESSAGE IN PHASE

The MESSAGE IN phase allows the target to request that messages be sent to the initiator from the target.

The target shall assert C/D, I/O, and MSG during the REQ/ACK handshake(s) of this phase.

5.1.9.2 MESSAGE OUT PHASE

The MESSAGE OUT phase allows the target to request that a message be sent from the initiator to the target. The target may invoke this phase at its convenience in response to the ATTENTION condition (see 5.2.1) created by the initiator.

The target shall assert C/D and MSG and negate I/O during the REQ/ACK handshake(s) of this phase. The target shall handshake byte(s) in this phase until ATN goes false.

If the target detects one or more parity error(s) on the message byte(s) received, it may indicate its desire to retry the message by asserting REQ after detecting ATN has gone false and prior to changing to any other phase. The initiator, upon detecting this condition, shall resend all of the previous message byte(s) sent during this phase. When resending more than one message byte, the initiator shall assert ATN prior to asserting ACK on the first byte and shall maintain ATN asserted until the last byte is sent as described in 5.2.1.

If the target receives all of the message byte(s) successfully (i.e., no parity errors), it shall indicate that it does not wish to retry by changing to any information transfer phase other than the MESSAGE OUT phase and transfer at least one byte. The target may also indicate that it has successfully received the message byte(s) by changing to the BUS FREE phase (e.g. ABORT or BUS DEVICE RESET messages).

5.1.10 SIGNAL RESTRICTIONS BETWEEN PHASES

When the SCSI bus is between two information transfer phases, the following restrictions shall apply to the SCSI bus signals:

- 1) The BSY, SEL, REQ, and ACK signals shall not change.
- 2) The C/D, I/O, MSG, and DATA BUS signals may change. When switching the DATA BUS direction from out to in (target to initiator), the target shall delay driving the DATA BUS by at least a data release delay plus a bus settle delay after asserting the I/O signal and the initiator shall release the DATA BUS no later than a data release delay after the transition of the I/O signal to true. When switching the DATA BUS direction from in to out (initiator to target), the target shall release the DATA BUS no later than a deskew delay after negating the I/O signal.
- 3) The ATN and RST signals may change as defined under the descriptions for the ATTENTION condition (5.2.1) and RESET condition (5.2.2).

5.2 SCSI BUS CONDITIONS

The SCSI bus has two asynchronous conditions: the ATTENTION condition and the RESET condition. These conditions cause the SCSI device to perform certain actions and can alter the phase sequence.

5.2.1 ATTENTION CONDITION

The ATTENTION condition allows an initiator to inform a target that the initiator has a message ready. The target may get this message at its convenience by performing a MESSAGE OUT phase.

The initiator creates the ATTENTION condition by asserting ATN at any time except during the ARBITRATION or BUS FREE phases.

The target may respond with the MESSAGE OUT phase.

The initiator shall keep ATN asserted if more than one byte is to be transferred. The initiator may negate the ATN signal at any time except it shall not negate the ATN signal while the ACK signal is asserted during a MESSAGE OUT phase. Normally, the initiator negates ATN while REQ is true and ACK is false during the last REQ/ACK handshake of the MESSAGE OUT phase.

5.2.2 RESET CONDITION

The RESET condition is used to immediately clear all SCSI devices from the bus. This condition shall take precedence over all other phases and conditions. Any SCSI device may create the RESET condition by asserting RST for a minimum of a reset hold time.

NOTE: The Q200 will never assert RST.

During the RESET condition, the state of all SCSI bus signals other than RST is not defined.

The Q200 will release all SCSI bus signals (except RST) within a bus clear delay of the transition of RST to true. The BUS FREE phase always follows the RESET condition.

The Q200 Series disk drives implement both a "hard" and a "warm" reset. The selection is performed by the "RO" Reset Option shorting plug on the drive unit. If "hard" reset is selected (the "RO" shorting plug has a jumper and is shorted) and the RST signal is detected, the drive will immediately reset all drive functions, including write, perform a complete servo-resynchronization, and go to the SCSI BUS FREE phase. This typically requires 6 or more

seconds to complete. If "warm" reset is selected (a jumper is not placed on the "RO" shorting plug) and the RST signal is detected, the drive will complete writing the current sector, perform a partial servo resynchronization and go to the SCSI BUS FREE phase. This typically requires 200 ms to complete.

5.3 SCSI BUS PHASE SEQUENCES

The order in which phases are used on the SCSI bus follows a prescribed sequence.

In all systems, the RESET condition can abort any phase and is always followed by the BUS FREE phase. Also, any other phase can be followed by the BUS FREE phase.

5.3.1 NONARBITRATING SYSTEMS

In systems where the ARBITRATION phase is not implemented, the allowable sequences shall be as shown in Figure 5-1. The normal progression is from the BUS FREE phase to SELECTION, and from SELECTION to one or more of the information transfer phases (COMMAND, DATA, STATUS, or MESSAGE).

5.3.2 ARBITRATING SYSTEMS

In systems where ARBITRATION phase is implemented, the allowable sequences shall be as shown in Figure 5-2. The normal progression is from the BUS FREE phase to ARBITRATION, from ARBITRATION to SELECTION or RESELECTION, and from SELECTION or RESELECTION to one or more of the information transfer phases (COMMAND, DATA, STATUS, or MESSAGE).

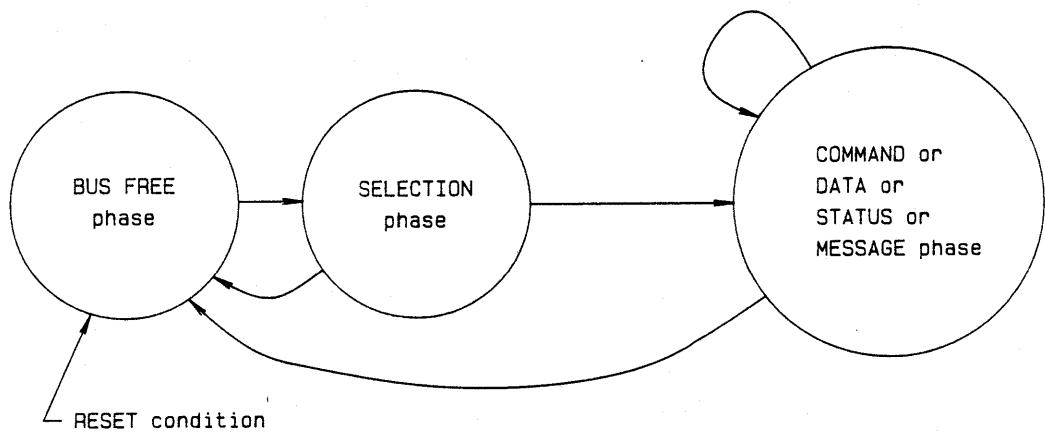


FIGURE 5-1. Phase Sequences without Arbitration

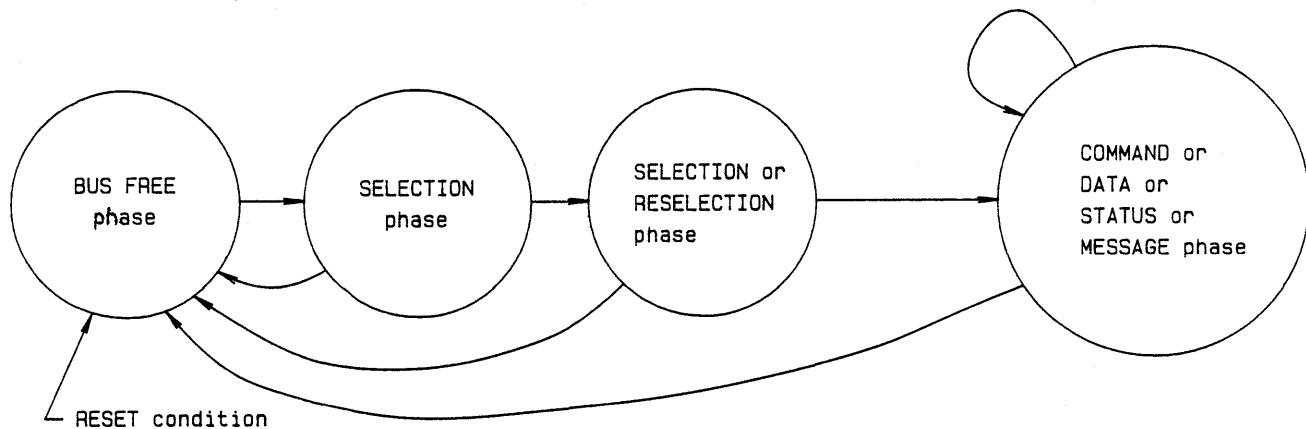


FIGURE 5-2. Phase Sequences with Arbitration

5.3.3 ALL SYSTEMS

There are no restrictions on the sequences between information transfer phases. A phase type may even be followed by the same phase type (e.g., a data phase may be followed by another data phase).

5.4 SCSI POINTERS

Consider the system shown in Figure 5-3 in which an initiator and target communicate on the SCSI bus in order to execute a command.

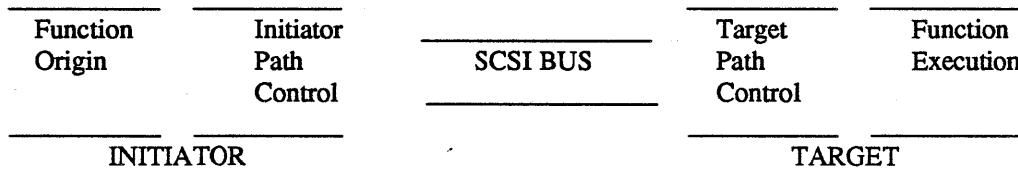


FIGURE 5-3. Simplified SCSI System

The SCSI architecture provides for two sets of three pointers within each initiator. The pointers reside in the initiator path control. The first set of pointers are known as the current (or active) pointers. These pointers are used to represent the state of the interface and point to the next command, data, or status byte to be transferred between the initiator's memory and the target. There is only one set of current pointers in each initiator. The current pointers are used by the target currently connected to the initiator.

The second set of pointers are known as the saved pointers. There is one set of saved pointers for each command that is currently active (whether or not it is currently connected). The saved command pointer always points to the start of the command descriptor block (see 6.1) for the current command. The saved status pointer always points to the start of the status area for the current command. At the beginning of each command, the saved data pointer points to the start of the data area. It remains at this value until the target sends a SAVE DATA POINTER message (see 5.5.2) to the initiator. In response to this message, the initiator stores the value of the current data pointer into the saved data pointer. The target may restore the current pointers to their saved values by sending a RESTORE POINTERS message (see 5.5.2) to the initiator. The initiator moves the saved value of each pointer into the corresponding current pointer. Whenever an SCSI device disconnects from the bus, only the saved pointer values are retained. The current pointer values are restored from the saved values upon the next reconnection.

5.5 MESSAGE SYSTEM SPECIFICATION

The message system allows communication between an initiator and target for the purpose of physical data path management.

5.5.1 MESSAGE PROTOCOL

The Q200 requires any SCSI device it communicates with to be able to receive and understand the COMMAND COMPLETE message. This is the only message an initiator communicating with the Q200 is required to implement. To maximize flexibility, however, the Q200 supports all single byte (non-extended) messages defined in the ANSI SCSI Specification. The Q200 will not send any message other than COMMAND COMPLETE until it is informed by the initiator that it can support other messages. The initiator does this in the SELECTION phase by asserting ATN prior to the SCSI bus condition of SEL true and BSY false. This informs the Q200 that the initiator can at least generate a MESSAGE REJECT message if it does not understand a message from the Q200.

If the Q200 detects the ATN signal asserted during SELECTION phase, it will enter MESSAGE OUT phase and take a message byte from the initiator. The Q200 expects an IDENTIFY message at this time. The NO OPERATION, BUS DEVICE RESET and ABORT messages will also function properly here. All other messages will be rejected with a MESSAGE REJECT message.

The IDENTIFY message the initiator sends to the Q200 should specify LUN=0. If the initiator supports disconnect/reconnect, it should set bit 6 in the IDENTIFY message (see Section 5.5.2). Unless bit 6 is set, the Q200 will not attempt to disconnect at any point during a command. Bit 6 should not be set unless the initiator can support receiving the SAVE DATA POINTER, RESTORE POINTERS, and DISCONNECT messages. Upon receipt of a legitimate IDENTIFY message, the Q200 will enter COMMAND phase and begin requesting command bytes from the initiator.

The Q200 supports the following messages:

Table 5-2
MESSAGE CODES SUPPORTED BY THE Q200 Series

Code	Description	Directions	Supported
00 _H	COMMAND COMPLETE	In (C)	
02 _H	SAVE DATA POINTER	In (C)	
03 _H	RESTORE POINTERS	In (C)	
04 _H	DISCONNECT	In (C)	
05 _H	INITIATOR DETECTED ERROR		Out (C)
06 _H	ABORT		Out
07 _H	MESSAGE REJECT	In (C)	Out (C)
08 _H	NO OPERATION		Out (C)
09 _H	MESSAGE PARITY ERROR		Out (C)
0A _H	LINKED COMMAND COMPLETE	In	
0B _H	LINKED COMMAND COMPLETE (WITH FLAG)	In	
0C _H	BUS DEVICE RESET		Out
80 _H - FF _H	IDENTIFY	In (C)	Out (C)

KEY: In = Q200 will originate message and send to initiator

Out = Q200 will support receiving message from initiator

(C) = Q200 will originate or support receiving of this message as an initiator during the COPY command (see 6.5.13).

5.5.2 MESSAGES

The single byte messages supported by the Q200 (Table 5-2) are listed along with their code values and their definitions.

COMMAND COMPLETE 00_H - This message is sent from the Q200 to an initiator to indicate that the execution of a command (or series of linked commands) has terminated and that valid status has been sent to the initiator. After successfully sending this message, the Q200 shall go into the BUS FREE phase by releasing BSY. This message must be supported by any initiator selecting the Q200.

NOTE: The command may have been executed successfully or unsuccessfully as indicated in the status.

SAVE DATA POINTER 02_H - This message is sent from the Q200 to direct the initiator to save a copy of the present active data pointer for the drive (see 5.4 for a definition of pointers). This message is usually sent just prior to the DISCONNECT message when disconnecting in the DATA phase.

NOTE: The Q200 drives with DisCacheTM will not send the SAVE DATA POINTER message prior to the first DISCONNECT message.

RESTORE POINTERS 03_H - This message is sent from the Q200 to direct the initiator to restore the most recently saved pointers to the active state. Saved pointers to the command, data, and status locations for the logical unit shall be restored to the active pointers. Command and status pointers shall be restored to the beginning of the present command and status areas. The data pointer shall be restored to the value at the beginning of the data area in the absence of a SAVE DATA POINTER message or to the value at the point at which the last SAVE DATA POINTER message occurred.

DISCONNECT 04_H - This message is sent from the Q200 to inform an initiator that the present physical data path is going to be broken (the Q200 plans to disconnect by releasing BSY), and that a later reconnect will be required in order to complete the current operation. If the initiator detects the BUS FREE phase (other than as a result of an ABORT or BUS DEVICE RESET message or a RESET condition) without first receiving a DISCONNECT or COMMAND COMPLETE message, the initiator shall consider this as a catastrophic error condition. This message shall not cause the initiator to save the data pointer.

INITIATOR DETECTED ERROR 05_H - This message is sent from an initiator to inform the Q200 that an error (e.g., parity error on the SCSI bus) has occurred that does not preclude a retry of the operation. Although present pointer integrity is not assured, a RESTORE POINTERS message or a disconnect followed by a reconnect, shall cause the pointers to be restored to their defined prior saved state. Depending on which command is being processed, the Q200 will either attempt to retransmit the data or it will terminate the command. If the data is to be retransmitted, the Q200 will first send a RESTORE POINTERS message and then continue. If the command is to be terminated a CHECK CONDITION status will result. The sense key shall be set to ABORTED COMMAND with a sense code of Initiator Detected Error.

ABORT 06_H - This message is sent from the initiator to the Q200 to clear the present operation. All pending data and status for the issuing initiator shall be cleared, and the Q200 shall go to the BUS FREE phase. Pending data and status for other initiators shall not be cleared. No status or ending message shall be sent for the operation.

MESSAGE REJECT 07_H - This message will be sent by the Q200 to an initiator when the last message received from that initiator was inappropriate or has not been implemented. This message will be interpreted in the same way when received from an initiator.

In order to indicate its intentions of sending this message, the initiator shall assert the ATN signal prior to its release of ACK for the REQ/ACK handshake of the message that is to be rejected. When a target sends this message, it shall change to MESSAGE IN phase and send this message prior to requesting additional message bytes

from the initiator. This provides an interlock so that the initiator can determine which message is rejected. The initiator must support this message if it asserts ATN during the initial SELECTION phase of the command.

NO OPERATION 08H - This message is sent, from an initiator in response to the Q200 request for a message, when the initiator does not currently have any other valid message to send. No action will be taken by the Q200 upon receipt of this message.

MESSAGE PARITY ERROR 09H - This message is sent from the initiator to the Q200 to indicate that one or more bytes in the last message it received has a parity error.

In order to indicate its intentions of sending this message, the initiator shall assert the ATN signal prior to its release of ACK for the REQ/ACK handshake of the message that has the parity error. This provides an interlock so that the target can determine which message has the parity error.

LINKED COMMAND COMPLETE 0AH - This message is sent from the Q200 to an initiator to indicate that the execution of a linked command has been completed and status has been sent. The initiator shall then set the pointers to the initial state for the next linked command. Refer also to Sections 6.1.7 and 6.1.8 below.

LINKED COMMAND COMPLETE (WITH FLAG) 0BH - This version of the LINKED COMMAND COMPLETE message is used when the flag bit of the command is set to one. Refer also to Sections 6.1.8 and 6.1.9 below.

BUS DEVICE RESET 0CH - This message is sent from an initiator to direct the Q200 to clear all current commands. This message forces the Q200 to undergo a warm reset (regardless of whether the reset option is set or not).

IDENTIFY 80H - FFH - These messages may be either sent from or received by the Q200 to establish the physical path connection between an initiator and the Q200.

Bit 7 - This bit is always set to one to distinguish an IDENTIFY message from the other messages.

Bit 6 - This bit is set to one only by the initiator. When set to one, it indicates that the initiator has the ability to accommodate disconnection and reconnection. Also, it indicates the SAVE DATA POINTER, RESTORE POINTERS and DISCONNECT messages are supported by the initiator.

Bits 5-3 - Reserved; must be set to zero.

Bits 2-0 - These bits specify a logical unit number (LUN) in a target. The only valid LUN on the Q200 is zero.

When an IDENTIFY message is sent from a target to an initiator during reconnection, an implied RESTORE POINTERS message shall be performed by the initiator prior to completion of this message.

5.5.3 MESSAGE ERROR RECOVERY STRATEGY

5.5.3.1 PARITY ERRORS DURING MESSAGE PHASE

Parity is an SCSI option. If the parity option is not enabled, the discussion in this section is not relevant. If the initiator detects a parity violation as the Q200 transfers a message byte to it, the initiator should ignore that message and immediately assert ATN (before releasing ACK on the offending byte). The initiator should continue receiving and ignoring bytes until the Q200 requests a MESSAGE OUT byte. At that time, the initiator should send a MESSAGE PARITY ERROR message. The Q200 will then cause the previous message to be retransmitted. If the retransmitted message is again garbled, the initiator should resend the MESSAGE PARITY ERROR message. Upon having the same message not communicated twice, the Q200 will give up, abort the command and get off the

bus. The sense key will be set to ABORTED COMMAND (0B_H) and the sense code set to Parity error on the SCSI bus (47_H).

If the Q200 detects a parity error during a MESSAGE OUT phase, it will continue to request MESSAGE OUT bytes until the initiator deasserts ATN. This procedure should guarantee that the initiator transfers the entire message containing the parity error. The drive then requests that the initiator resend the message by not changing the phase and requesting more MESSAGE OUT bytes. If the second attempt still does not result in a good message, the current command will be terminated, the status set to CHECK CONDITION, the sense key set to ABORTED COMMAND (0B_H) and the sense code set to Parity error on the SCSI bus (47_H).

5.5.3.2 REJECTED MESSAGES

The Q200 will send MESSAGE REJECT to the initiator any time it receives a message it does not support (e.g., extended messages) or a message that does not make sense at that time (e.g., the initiator sending COMMAND COMPLETE).

When the Q200 receives MESSAGE REJECT from the initiator, it attempts to resend the message. If the message is again rejected, it takes action based on which command was rejected.

COMMAND COMPLETE: Q200 goes to BUS FREE. This is not considered an error.

DISCONNECT: The Q200 will stop trying to disconnect at this time. It may attempt to disconnect again at a later time.

IDENTIFY: The Q200 only sends this when attempting to reconnect. Rejecting this message during reconnect will terminate the present command. A CHECK CONDITION status will be posted, the sense key set to HARDWARE ERROR (4_H) and the sense code set to Message reject error (43_H). The Q200 then goes immediately to BUS FREE.

LINKED COMMAND COMPLETE: The Q200 will go immediately to BUS FREE. The sense key will be set to HARDWARE ERROR (4_H) and the sense code set to Message reject error (43_H).

LINKED COMMAND COMPLETE WITH FLAG: Same procedure as for LINKED COMMAND COMPLETE.

MESSAGE REJECT: The Q200 will immediately terminate the current command with a CHECK CONDITION status. The sense key will be set to HARDWARE ERROR (4_H) and the sense code set to Message reject error (43_H).

RESTORE POINTERS: This message is only sent by the Q200 when attempting to retry a command or part of a command. The Q200 will abort the retry process. The reject of the RESTORE POINTERS message in itself will not cause an error condition.

SAVE DATA POINTERS: Same procedure as for DISCONNECT message.

6.0 COMMAND SPECIFICATION

This section defines the SCSI commands supported by the Q200 Intelligent Disk Drives. Each command consists of a group of command bytes and an optional associated group of data bytes. After ARBITRATION, SELECTION, and any optional initial messages, the Q200 will enter the COMMAND phase and request the command bytes. The opcode of the command will determine how many bytes the Q200 will request. After receipt of the command bytes, the Q200 will enter DATA phase if a data structure is defined for the command. The data structure in a single command may transfer one or more logical blocks of data. The Q200 may disconnect from the SCSI bus to allow activity by other SCSI devices while it is preparing to transfer data.

Upon command completion (successful or unsuccessful), the Q200 returns a status byte to the initiator. Since most error and exception conditions cannot adequately be described with a single status byte, one status code, CHECK CONDITION, indicates that additional information is available. The initiator may issue a REQUEST SENSE command to retrieve this additional information.

The subsections on each command will describe the sequence of bytes an initiator must transfer to the drive during the COMMAND phase in order to request a desired action from the drive. The subsections on each command may further describe the format of data bytes that may be transferred to or from the initiator following the COMMAND phase, depending on the function of the command.

6.1 COMMAND DESCRIPTOR BLOCK (CDB)

A request to a device is made by transferring a Command Descriptor Block (CDB) to the target device. The Q200 supports two basic command formats, six-byte commands and ten-byte commands. The ANSI SCSI Specification refers to the six-byte commands as belonging to Group 0 and the ten-byte commands as belonging to Group 1. The command format is recognizable from the opcode of the command. Six-byte commands have opcodes in the range of 00_H to $1F_H$; ten-byte commands have opcodes in the range of 20_H to $3F_H$. This section defines the CDB fields, whose uses in particular commands are fully described in Section 6.5.

6.1.1 OPCODE

The opcodes (operation codes) for Groups 0 and 1 are listed in Table 6-1.

Table 6-1
COMMANDS SUPPORTED BY THE Q200 Series Drives

OPCODE	COMMAND	SECTION
00 _H	TEST UNIT READY	6.5.1
01 _H	REZERO UNIT	6.5.2
03 _H	REQUEST SENSE	6.5.3
04 _H	FORMAT UNIT	6.5.4
07 _H	REASSIGN BLOCKS	6.5.5
08 _H	READ	6.5.6
0A _H	WRITE	6.5.7
B _H	SEEK	6.5.8
12 _H	INQUIRY	6.5.9
15 _H	MODE SELECT	6.5.10
16 _H	RESERVE	6.5.11
17 _H	RELEASE	6.5.12
18 _H	COPY	6.5.13
1A _H	MODE SENSE	6.5.14
1B _H	START/STOP UNIT	6.5.15
1D _H	SEND DIAGNOSTIC	6.5.16
25 _H	READ CAPACITY	6.5.17
28 _H	READ EXTENDED	6.5.18
2A _H	WRITE EXTENDED	6.5.19
2B _H	SEEK EXTENDED	6.5.20
2F _H	VERIFY	6.5.21
37 _H	READ DEFECT DATA	6.5.22
3B _H	WRITE BUFFER	6.5.23
3C _H	READ BUFFER	6.5.24

6.1.2 LOGICAL UNIT NUMBER (LUN):

The Logical Unit Number for the Q200 is zero. LUNs with the value of 1-7 are not accepted. The LUN is found in two places, bits 5-7 of byte 1 of the Command Descriptor Block and bits 0-2 of the IDENTIFY message field (refer to Section 5.5.2). The LUN field in both the CDB and IDENTIFY message field must be zero. If a LUN other than zero is specified, the command (any command except INQUIRY and REQUEST SENSE) will be rejected with a CHECK CONDITION status with the sense key set to ILLEGAL REQUEST - Invalid LUN.

6.1.3 LOGICAL BLOCK ADDRESS (LBA):

The 21-bit Logical Block Address signifies the first or starting block of an operation. (Each Logical Block can be 512, 1024, or 2048 bytes in length.) The MODE SELECT command is used to configure the Block Length. Each Q200 is configured with a Logical Block Length of 512 bytes.

6.1.4 TRANSFER OR ALLOCATION LENGTH:

The transfer length in blocks, or the allocation length in bytes, is the number of blocks or bytes to be transferred for a single command.

For Group 0 commands, acceptable values are 00-FF_H. For block lengths, a value of 00_H transfers 256 blocks. See allocation byte lengths under specific commands.

For Group 1 commands, acceptable values are 0000-FFFFH. A value of 0000H transfers no data blocks.

6.1.5 VENDOR UNIQUE (VU):

Vendor unique bits allow for slight command modifications or special conditions to be selected. Vendor unique bits are to be filled with zero.

6.1.6 RESERVED (R):

These bits are reserved and must be set to zero by the initiator.

6.1.7 FLAG BIT (F):

If the link bit is zero, then the flag bit must be zero. If the link bit is one when the command terminates, the target sends the LINKED COMMAND COMPLETE message if the flag bit is zero, or the LINKED COMMAND COMPLETE (WITH FLAG) message if the flag bit is one.

6.1.8 LINK BIT (L):

If the link bit is one when the command terminates, the target sends one of the two messages defined by the flag bit (above). If the link bit is zero, the target sends the normal COMMAND COMPLETE message.

6.1.9 RELATIVE ADDRESS (RELADR) BIT:

This bit is found only in Group 1 commands (opcodes 20H - 3FH). This bit is set to one to indicate that the logical block address portion of the command descriptor block is a displacement. This positive or negative (two's complement) displacement is added to the logical block address last accessed on the logical unit to form the logical block address for this command. A previous command of the linked group, including the current command, must have accessed a block of data on the logical unit. Use of this bit is not supported by the Q200 disk drives.

NOTE: If the reserved bits of the IDENTIFY message or the vendor unique, reserved, or relative address bits of the CDB are non-zero; then a CHECK CONDITION status is returned with the sense key set to ILLEGAL REQUEST and sense code set to Illegal use of bit or byte in CDB or IDENTIFY message.

6.2 COMPLETION STATUS BYTE

A status byte (Table 6-2) shall be sent from the target to the initiator during the STATUS phase at the termination of each command as specified in Tables 6-3 and 6-4 unless the command is cleared by an ABORT message, by a BUS DEVICE RESET message, or by a "hard" or "warm" RESET condition. Table 6-4 defines the status codes as implemented on the Q200.

Table 6-2
STATUS BYTE

Byte	Bit 7	6	5	4	3	2	1	0
0	RESERVED	VENDOR UNIQUE		STATUS BYTE CODE		VENDOR UNIQUE		

Table 6-3
STATUS BYTE CODE BIT VALUES

Bits of Status Byte								Status(es) Represented	Q200 Series
7	6	5	4	3	2	1	0		
R	V	V	0	0	0	0	V	GOOD	Implemented
R	V	V	0	0	0	1	V	CHECK CONDITION	Implemented
R	V	V	0	0	1	0	V	CONDITION MET/GOOD	Not Used
R	V	V	0	0	1	1	V	Reserved	Not Used
R	V	V	0	1	0	0	V	BUSY	Implemented
R	V	V	0	1	0	1	V	Reserved	Not Used
R	V	V	0	1	1	0	V	Reserved	Not Used
R	V	V	0	1	1	1	V	Reserved	Not Used
R	V	V	1	0	0	0	V	INTERMEDIATE/GOOD	Implemented
R	V	V	1	0	0	1	V	Reserved	Not Used
R	V	V	1	0	1	0	V	INTERMEDIATE/CONDITION MET/GOOD	Not Used
R	V	V	1	0	1	1	V	Reserved	Not Used
R	V	V	1	1	0	0	V	RESERVATION CONFLICT	Implemented
R	V	V	1	1	0	1	V	Reserved	Not Used
R	V	V	1	1	1	0	V	Reserved	Not Used
R	V	V	1	1	1	1	V	Reserved	Not Used

KEY: R = Reserved Bit (zero)
V = Vendor Unique Bit (zero)

Table 6-4
Q200 Series STATUS CODES

STATUS BYTE CODE	MEANING
GOOD (00H)	COMMAND HAS COMPLETED WITH NO ERROR CONDITIONS.
INTERMEDIATE/ GOOD (10H)	LINKED COMMAND HAS COMPLETED WITH NO ERROR CONDITIONS. This code is issued for every command in a series of linked commands (except the last, for which the link bit is zero) unless CHECK CONDITION or RESERVATION CONFLICT is appropriate. If this code is not returned, the chain of linked commands is broken.
CHECK CONDITION (02H)	AN ABNORMAL CONDITION HAS OCCURRED. This code is issued whenever an error, exception, or abnormal condition occurred during the execution of a command (the command may or may not have completed successfully). A REQUEST SENSE command should be issued following a CHECK CONDITION status to determine the nature of the condition.
BUSY (08H)	TARGET IS BUSY. This code is issued if the target is busy executing a previous command, after a disconnection has occurred, or if the target is otherwise unable to accept a command. The command that received the BUSY status is discarded and no action is taken; it may be reissued as if it was issued for the first time.
RESERVATION CONFLICT (18H)	ATTEMPTED ACCESS OF A RESERVED BLOCK, OR DEVICE. This code is issued when a device attempts to access a reserved in which case the command will be rejected. See RESERVE command.

6.3 READING AND WRITING

6.3.1 LOGICAL BLOCKS

The basic functions of a disk drive are to store and retrieve information. The information on the Q200 Series disk drives is stored and retrieved as groups of 8-bit bytes. These groups of bytes have a standard size and are referred to as logical blocks. The size of a logical block, the number of bytes contained in each block, is determined by the user via the MODE SELECT command. The Q200 will support three possible logical block sizes: 512 bytes per logical block, 1024 bytes per block, and 2048 bytes per block. The drive is initially configured as 512 bytes per logical block. A subtle feature of the Q200 is that changing the block size will not destroy data on the drive. The changed block size will just change the grouping of the information (e.g., 1024-byte blocks would reconfigure as two 512-byte blocks).

In order to specify which logical block a user wants to store away or retrieve, each block has its own address. The address of each logical block is called the Logical Block Address and is frequently referred to as the LBA. Thus, to the user, the Q200 appears as a series of logical blocks, each one with its own unique address. The first logical block on the drive is always address 0. The address of the last logical block on the drive varies depending on the capacity of the drive and the size of the logical block.

6.3.2 LOGICAL BLOCK ADDRESS RANGES

	Block Size	Valid Logical Block Address Range	
		Decimal	Hex
Q250	512 or 200_H	0 - 103,697	0 - 19511_H
	1024 or 400_H	0 - 51,848	0 - $CA88_H$
	2048 or 800_H	0 - 25,923	0 - 6543_H
Q280	512 or 200_H	0 - 156,369	0 - $262D1_H$
	1024 or 400_H	0 - 78,184	0 - 13168_H
	2048 or 800_H	0 - 39,091	0 - $98B3_H$

If an initiator attempts to access a logical block at an address outside the valid address range, the command will not be executed. The drive will return a CHECK CONDITION status with the sense key set to ILLEGAL REQUEST - Invalid Logical Block Address was specified.

6.3.3 TRANSFERRING DATA

The actual storing away and retrieving of logical blocks is done with the WRITE and READ commands. WRITE will place data on the disk and READ will retrieve data from the disk. In both the READ and WRITE commands, the user (initiator) must specify the logical block address of the first logical block to be transferred and the total number of blocks to be transferred. The interpretation of a value of zero in the number of blocks to be transferred depends on which form of the READ or WRITE command is used. If the long or extended form of the commands is used (opcodes 28_H and $2A_H$), then a transfer length of zero implies transfer no data. In this case, the command behaves just like a SEEK command. If the shorter or non-extended form of the commands is used (opcodes 08_H and $0A_H$), then a transfer length of zero means transfer 256 blocks during the command. This is because the shorter command format has only one byte to specify transfer length.

The number of bytes transferred during the DATA phase following the command bytes is determined by the transfer length specified in the command and the logical block size

$$\text{Number of Data Bytes} = \text{Logical Block Size (in bytes/block)} * \text{Transfer Length (in blocks)}$$

Note that in case of errors, fewer number of bytes will sometimes be transferred.

The rate at which data is transferred to or from the disk depends solely on how densely the data is written on the disk and how fast the disk is rotating. Neither of these quantities is readily changeable, so the disk transfer rate should be regarded as a constant. The transfer rate when the head is actively over data is 10 Mbit/sec (1.25 Mbyte/sec). The sustained disk transfer rate is slower because interspersed among the data are servo information and intersector gaps which absorb about 20% of the disk space so the sustained disk transfer rate is about 8 Mbit/sec (1.0 Mbyte/sec). If the transfer rate across the SCSI bus is slower than this sustained disk transfer rate, the drive will either run out of data or begin to have data pile up in its buffer (depending on whether the transfer is a WRITE or READ, respectively). Other disk drive subsystems might require the user to interleave sectors on the disk to match the sustained disk transfer rate with the SCSI bus transfer rate. For more flexibility, the Q200 includes a large on-board buffer (14K bytes). All READ and WRITE transfers are channeled through this buffer which allows the drive to transfer steadily at transfer rates slower than 1.0 Mbyte/sec.

Some Q200 Series drives feature DisCache, an on-board 60K-byte look-ahead disk cache. All READ and WRITE data transfers are channeled through the cache which also serves as a data transfer buffer. The user can, by use of the MODE SELECT command, enable/disable the cache function as well as configure various cache parameters to best fit the application. Please refer to the section on the MODE SELECT command and Appendix B for a more detailed discussion of the caching function.

6.4 CONFIGURING THE Q200 OPERATING MODE

6.4.1 Q200 OPERATING MODES

The manner in which the Q200 operates can be changed in order to more precisely match its features with the user's requirements. This is referred to as changing the Operating Mode of the drive and is done through the use of the MODE SELECT and MODE SENSE commands. The MODE SELECT command is used to set the Operating Mode of the Q200 in order to achieve the desired behavior. The MODE SENSE command may be used to determine the current Operating Mode of the drive, or it may be used to gather other information that would be useful in setting the Operating Mode.

An example of a common user adjusted Operating Mode parameter is the number of retries the drive should attempt when it detects a read error. If an initiator wished to retry only once when the drive detected incorrect data, the initiator's software could first issue a MODE SENSE command to the Q200 to determine how the drive is currently configured. The data returned to the initiator during the DATA phase of the MODE SENSE command would contain the current Mode Parameters, including how many retries the drive will attempt when it detects its read data incorrectly. If the number of retries is not set to the desired number, the initiator could then send a MODE SELECT command with the data byte corresponding to the Retry Count Parameter set to one. From then on, READ commands issued from that initiator or any another initiator would only retry once to try to get good data.

6.4.2 OPERATING MODE TABLES

The Q200 Operating Mode is determined by the contents of the Current Mode Table. Because of the need for frequent and fast access to its parameters, the Current Mode Table is stored in RAM. This table is referenced whenever there is more than one possible way for the drive to behave. In the previous example, the initiator will have set the Retry Count in the Current Mode Table to one. Each time a READ command is executed, the firmware checks the Current Mode Table and looks up the Retry Count Parameter. It then uses this number to determine how many times it should reattempt to read a block if the data is detected to be incorrect.

Storing the Current Mode Table in RAM causes the drive's Operating Mode to be lost when power is removed or the drive is reset. For this reason, a copy of the Current Mode Table can be saved on the disk itself. This saved copy is called the Saved Mode Table. The Saved Mode Table is stored on a hidden area of the disk that is not accessible by the user. This way, no user data storage area is sacrificed nor can the table be inadvertently wiped out by the user. The Saved Mode Table is used to load the Current Mode Table at power-up or whenever the drive is reset. The contents of the Current Mode Table can be copied onto the Saved Mode Table through the use of the MODE SELECT command. If the Save Parameters bit is set to one in the MODE SELECT command (bit 0 of byte 1) then the entire contents of the Current Mode Table will be transferred to the Saved Mode Table when the command is completed. Note that if the MODE SELECT command makes changes to the Current Mode Table, these changes will be implemented first and then incorporated in the Saved Mode Table. (In other words, if a MODE SELECT command contains new mode information and the Save Parameters bit is set, the new mode information will appear in the Saved Mode Table as well as in the Current Mode Table.)

The Saved Mode Table is a very powerful tool for the user. Using the Save Parameters bit allows the drive to be permanently configured to the user's requirements with the execution of a single MODE SELECT command. This makes available the option of doing a once-in-a-lifetime configuration (e.g., at the user's factory) and eliminating the burden from the device driver of sending vendor specific MODE SELECT commands after every reset.

There are two other Mode Tables supported by the Q200 of interest to the user. These are the Changeable Parameters Table and the Default Mode Table. Both of these tables are contained in the firmware ROM and are for reference only. They are not modifiable by the user. These two tables can be referenced only by the MODE SENSE command.

The Changeable Parameters Table allows an initiator to determine which Operating Mode parameters the Q200 will allow to be changed via a MODE SELECT command. If a parameter in the Current Mode Table is changeable, the corresponding entry will have all bits set to one. Before trying to change any Operating Mode parameter with a MODE SELECT command, the software driver should first check to make sure that the device will support changing that parameter.

The Default Mode Table is provided to allow the user to restore the drive to a known or reference condition. If for any reason the Current and Saved Mode Tables are inaccessible, the Q200 will use the Default Mode Table.

6.5 DESCRIPTION OF COMMANDS

Following are descriptions of individual commands listed in Table 6-1. Refer to the REQUEST SENSE command (6.5.3) for a description of error codes and sense keys. Refer to the Command Descriptor Block section (6.1) for definitions of the CDB fields.

6.5.1 TEST UNIT READY COMMAND: OPCODE = 00 (hex)

Table 6-5
TEST UNIT READY COMMAND

Bit Byte	7	6	5	4	3	2	1	0
0								OPCODE = 00 _H
1			LUN = 0			RESERVED = 0		
2-4				RESERVED = 0				
5	VU = 0			RESERVED = 0			F	L

The TEST UNIT READY command (Table 6-5) provides a means to check if the drive is up to speed and ready to accept commands requiring disk access.

When the Q200 drive is ready, the completion status byte will indicate GOOD (00_H). If the drive is not ready, the completion status byte will indicate CHECK CONDITION (02_H), with the sense key set to NOT READY (02_H).

The commands which do not require disk access are:

REQUEST SENSE	INQUIRY *
RESERVE	SEND DIAGNOSTICS
RELEASE	MODE SELECT (if "SAVE PARAMETER BIT NOT SET")
START/STOP UNIT	MODE SENSE (if SAVED PARAMETERS ARE NOT REQUESTED)
WRITE BUFFER	
READ BUFFER	

* NOTE: If the Q200 is not ready when the INQUIRY command is received, the disk is not accessed and the default INQUIRY data is returned (see 6.5.9).

These commands will be executed after a delay of one second (two seconds for the Q200 with DisCache) maximum following application of power to the Q200. They will not return CHECK CONDITION status with the NOT READY sense key even if the drive is not ready. All other commands could require disk access. (Refer to START/STOP UNIT command, Section 6.5.15.1.)

6.5.2 REZERO UNIT COMMAND: OPCODE = 01 (hex)

Table 6-6
REZERO UNIT COMMAND

Byte	7	6	5	4	3	2	1	0
0	OPCODE = 01 _H							
1	LUN = 0			RESERVED = 0				
2-4	RESERVED = 0							
5	VU = 0		RESERVED = 0			F	L	

The REZERO UNIT command (Table 6-6) requests that the drive actuator be repositioned to cylinder zero and a recalibration procedure (which is automatically performed upon powering up and starting the motor) be performed.

6.5.3 REQUEST SENSE COMMAND: OPCODE = 03 (hex)

6.5.3.1 REQUEST SENSE COMMAND STRUCTURE:

Table 6-7
REQUEST SENSE COMMAND

Byte	7	6	5	4	3	2	1	0
0	OPCODE = 03 _H							
1	LUN = 0			RESERVED = 0				
2-3	RESERVED = 0							
4	ALLOCATION LENGTH (BYTES)						F	L
5	VU = 0		RESERVED = 0					

The REQUEST SENSE command (Table 6-7) requests that the sense data be sent to the initiator.

The sense data applies to the previous command from the same initiator. The sense data is saved for this initiator until retrieved by this initiator or until another command is received from this initiator. Commands from initiators will not alter the sense data saved for other initiators.

The allocation length specifies the maximum number of bytes that the initiator has allocated for returned sense data. The exception is a value of 00, which indicates that four bytes of data will be returned. The acceptable values are 00-FF_H. The target will terminate the DATA IN phase when the allocated number of bytes have been transferred or when all available sense data have been transferred to the initiator, whichever is less. For the Q200, an allocation length of 18 (12_H) is recommended.

If the allocation length prevents all available sense data from being sent to the initiator, the initiator may not recover the remaining portion with a subsequent REQUEST SENSE command.

The Q200 uses the extended format for sense data (Table 6-8).

6.5.3.2 EXTENDED SENSE FORMAT:

Table 6-8
EXTENDED SENSE FORMAT

Byte	Bit	7	6	5	4	3	2	1	0
0	VALID	ERROR CLASS AND CODE = 70 _H							
1		SEGMENT NUMBER = 0							
2	FM=0	EOM=1	ILI=0	R=0	SENSE KEY				
3 - 6	(MSB)	INFORMATION BYTE						(LSB)	
7		ADDITIONAL SENSE LENGTH (BYTES) = 0A _H							
8-11		RESERVED = 0							
12		ADDITIONAL SENSE CODE							
13		RESERVED = 0							
14		FRU FAILED = 0							
15	FPV	C/D	VU	VU	BPV	BIT POINTER = 0			
16-17	(MSB)	FIELD POINTER						(LSB)	

INFORMATION VALID (VALID): Set to 1 if the information bytes contain valid data.

ERROR CLASS AND CODE: Error class and code is set to 70_H to indicate extended status.

SEGMENT NUMBER: The segment number contains the number of the current segment descriptor, if the extended sense is in response to a COPY command. Otherwise, this field is set to zero for the Q200.

FILE MARK (FM): Indicates that a filemark has been read. This bit is set to zero for the Q200.

END-OF-MEDIUM (EOM): Indicates that an end-of-medium condition exists. This bit is set to zero for the Q200.

INCORRECT LENGTH INDICATOR (ILI): Indicates that the requested logical block length does not match that on the medium. This bit is set to zero for the Q200 except for the READ BUFFER command, when the requested transfer length is greater than the buffer size.

RESERVED (R): Must be zero for the Q200.

SENSE KEY: The values for sense key data are found in the Table 6-9.

INFORMATION BYTES:

If the VALID bit is 1, the information bytes are valid and are defined as follows:

1. The logical block address associated with the sense key, for commands other than COPY.
2. The difference of the requested number of blocks minus the actual number of blocks copied for the current segment descriptor of a COPY command.
3. The difference of the requested number of bytes minus the actual number of bytes transferred for the READ BUFFER or WRITE BUFFER command.

ADDITIONAL SENSE LENGTH: Additional sense length is the number of additional bytes to follow. This value will not be truncated if the allocation length in the command descriptor block is too small to transfer all of the additional sense bytes. The Q200 uses a value of $0A_H$ in byte 7) except in certain cases with the COPY command (see Section 6.5.13.5).

ADDITIONAL SENSE CODE: Refer to Table 6-9. (This field is not used in certain cases with the COPY commands - refer to Section 6.5.13.5.)

FIELD REPLACEABLE UNIT (FRU) FAILED: The format of this byte is vendor unique. A value of zero means that no FRU is to be reported.

FIELD POINTER VALID (FPV): Bit 7 set to zero indicates that the target does not implement the functions supported by C/D bit, and BVP bit, and bytes 16 and 17; therefore, these bits and fields are not valid. FPV bit set to one indicates that the field pointer bytes 16 and 17, the C/D bit, and BVP bits are significant.

CONTROL/DATA BIT (C/D): C/D bit set to one indicates that the value reported in the field pointer is the Command Descriptor Block (CDB) byte number for which an ILLEGAL REQUEST sense key was issued. C/D bit of zero indicates that the value reported in the field pointer is the byte number of the DATA phase for which an ILLEGAL REQUEST sense key was issued.

VENDOR UNIQUE: Bits 5 and 4 of byte 15 are vendor unique.

BIT POINTER VALID (BVP): BVP bit of zero indicates that the target does not implement this function; therefore, the bit pointer field is not valid. A BVP bit of one indicates that the bit pointer field (Bits 0 - 2) is valid.

FIELD POINTER: Bytes 16 and 17 of the extended sense data format represent an index into the Command Descriptor Block fields or the DATA phase fields for which an ILLEGAL REQUEST sense key was issued. In the case of multiple byte fields, the field pointer bytes index the Most Significant Byte of the field for which an ILLEGAL REQUEST sense key was issued. These bytes are only valid if the Field Pointer Valid (FPV) bit is set to one.

Table 6-9
SENSE ERROR CODES

Sense Key (hex)	Additional Sense Code (hex)	Description
0		NO SENSE. Indicates that there is no specific sense key information to be reported. This would be the case for a successful command or a command that received a CHECK CONDITION status because one of the FM, EOM, or ILI bits is set to one.
	00	No additional information.
1		RECOVERED ERROR. Indicates that the last command completed successfully with some recovery action performed by the Q200. Details may be determined by examining the additional sense bytes and the information bytes.
	00	No additional information.
	11	Data error with uncorrectable ECC was recovered with retries. Sector should be reassigned.
	15	Recoverable seek error.
	17	Data was recovered with retries. An error was recovered within the maximum number of retries set in MODE SELECT.
	18	Data was recovered using ECC. An error was recovered using the ECC, after using the maximum number of retries.
	32	Defect List is full. Indicates that the defect list is now full. NO more defects may be entered unless the drive is formatted using the FORMAT UNIT command (Opcode 04H). Since the FORMAT UNIT command typically rearranges disk blocks, data must be backed up first.
	98	Data was recovered using the ECC at final retry.
2		NOT READY. Indicates that the target cannot be accessed. Operator intervention may be required to correct this condition.
	00	No additional information.
	B0	Drive is in recal. Drive motor is up to speed, but the drive is in the process of recalibration. Recalibration includes synchronizing the timing controller chip up to edges, seeking to logical cylinder 822 and setting AGC (Automatic Gain Control) values for each head, and seeking to logical track 0 and setting AGC values.
	B1	Drive is spinning up. Drive motor has not yet reached or maintained operating speed for at least 2 seconds.

B2 Drive is waiting for START/STOP UNIT command. Indicates that the drive is waiting for the START/STOP UNIT command (Opcode 1B_H) to start the drive before spinning up and recalcing.

3 MEDIUM ERROR. Indicates that the command terminated with a non-recovered error condition that was probably caused by a flaw in the medium or an error in the recorded data.

00 No additional information.

10 Repeatable qual errors. The timing controller chip detected an error (unrecoverable after using 32 retries) in the sector and echo pulse timing relationship (in a servo wedge). Once a qual error is detected, writing is inhibited in the sector following the wedge in which the error occurred.

11 Uncorrectable data error. An error of greater than 17 bits has been detected and is unrecoverable after using enough retries to get a stable syndrome, or using the maximum number of retries, whichever occurs first. No data will be transferred to the host unless the TB bit in MODE SELECT is set to one.

13 No sync field was found when trying to read data. No sync pattern was detected in the data field within the allotted window. Data was unreadable after using the maximum number of retries.

14 An illegal sync field was found when trying to read data. An invalid sync pattern occurred in the data field and was unrecoverable after using the maximum number of retries. No data will be transferred to the host.

1C No defect lists (P & G) found. After power up, disk cannot read the defect map into buffer RAM.

32 Running out of spares. There are no more spare sectors on the disk to be used as replacements for defective sectors.

91 Uncorrectable data error at final retry.

93 No sync field was found when trying to read data at final retry.

94 An illegal sync field was found when trying to read data at final retry.

4 HARDWARE ERROR. Indicates that the target detected a non-recoverable hardware failure while performing the command.

00 No additional information.

02 Failed recal during seek.

03 Write fault error. A write unsafe condition was detected and the WRITE command was aborted. No retries are attempted by the drive and the WRITE command must be reissued to try again.

15 Seek error.

- 19 Bad defect list. Reference defect map in buffer RAM was found corrupted.
- 31 Fail to read/write defect lists in FORMAT command. Indicates a failure in storing the defect list.
- 32 Defects overflow. The defects have overflowed the defect map in the buffer RAM.
- 43 Message reject error. The target has rejected either the IDENTIFY, COMMAND COMPLETE or the REJECT MESSAGE message.
- 44 SCSI hardware/firmware error.
- 45 Select/reselect failure. The initiator has not responded to the reselection attempt within the selection timeout delay, and the reselection has been aborted.
- A0 DICEY miscellaneous hardware error. A miscellaneous error (other than 4A1, 4A2, 4A3, 403 or 4AE) has been detected in the data controller chip.
- A1 DICEY data collision. The data controller IC received data pulses which were too high in frequency from the data synchronizer chip while attempting to read data.
- A2 DICEY servo collision. The data controller IC received servo data pulses which were too high in frequency from the data synchronizer chip while attempting to read servos.
- A3 DICEY attempting to write over servos. The data controller IC was attempting to write during a servo wedge. NO retries are attempted and all writing is aborted.
- A4 DICEY failing to abort READ/WRITE when an ABORT is issued.
- A5 DICEY found with an incorrect pointer. The data controller IC has detected a mismatch between the disk channel pointer and the SCSI channel pointer after a READ or WRITE.
- A6 DICEY aborted READ/WRITE for no apparent reason. The data controller IC inhibited a READ or WRITE without an ABORT being issued.
- A8 DICEY indicates an impossible condition.
- A9 RAM parity failure when not accessed.
- AE DICEY timed out between wedges (FYLO error). The data controller IC did not receive a wedge signal from the timing controller chip within the timeout period.
- AF Miscellaneous FYLO hardware error (in the timing controller chip).
- 5 ILLEGAL REQUEST. Indicates that there was an illegal parameter in the command descriptor block or in the additional parameters supplied as data for some commands. If the Q200 detects an invalid parameter in the CDB, then it will terminate the command without altering the medium. If the Q200 detects an invalid parameter in the additional parameters supplied as data, then it may have already altered the medium.
- 00 No additional information.
- 19 Defect list is not in order. All defects must be arranged in ascending order before the defect list will be accepted.

20 Invalid command operation code. The host has issued a command which the Q200 does not support or cannot recognize.

21 Invalid logical block address was specified. The maximum LBAs are:

	Q250	Q280
512 bytes/block	103,697	156,369
1024 bytes/block	51,848	78,184
2048 bytes/block	25,923	39,091

24 Illegal use of bit or byte in CDB or IDENTIFY message (if one was sent). Invalid field in the CDB or IDENTIFY message.

25 Invalid LUN. A logical unit number other than zero was used in either the CDB or the IDENTIFY message (if one was sent). The Q200 supports only LUN zero.

26 Invalid field in parameter list.

90 Invalid transfer length.

A0 Couldn't read or write saved MODE SELECT parameters. No copy of the mode parameters could be found on the disk. The mode parameters must first be saved before they can be read.

A1 Disconnect bit wasn't set for COPY command. This bit must be set for COPY to take place.

6 UNIT ATTENTION. MODE SELECT parameters have been changed by another initiator, or the Q200 has been reset by one of the following:

- 1) BUS DEVICE RESET message
- 2) "Hard" RESET condition (RST asserted)
- 3) Power on RESET

The UNIT ATTENTION condition will persist until receipt of any command other than INQUIRY or REQUEST SENSE; such command will not be performed and will report the CHECK CONDITION status, allowing the initiator to discover and clear the UNIT ATTENTION condition by issuing a REQUEST SENSE command. An INQUIRY command will be performed with no effect to the UNIT ATTENTION condition. A REQUEST SENSE command received before the target reports CHECK CONDITION status for the UNIT ATTENTION condition will cause the condition to be reported and cleared.

00 No additional information.

29 RESET occurred. Indicates a power up RESET, a bus RESET, or a BUS DEVICE RESET has occurred. MODE SELECT parameters revert to the saved parameters stored on the disk.

2A MODE SELECT parameters have been changed. They have been changed via the MODE SELECT command by another initiator.

A COPY ABORTED. Indicates a COPY command was aborted due to an error condition on the source device, the destination device, or both.

00 No additional information

45 COPY aborted due to failed target selection. While attempting a COPY command, the target device could not be selected.

A0 COPY aborted due to target error. The source or destination device detected some error during execution of the COPY command.

B ABORTED COMMAND. Indicates that the Q200 aborted the command. The initiator may be able to recover by trying the command again.

00 No additional information.

40 Buffer RAM parity error. A parity error was detected when reading from the buffer RAM.

43 Message reject error.

47 Parity error on the SCSI bus.

48 Initiator detected error. If the target receives this error, it may retry the previous operation using the following sequence:

1. Change the phase to MESSAGE IN, and send to the initiator the RESTORE POINTERS message.
2. Return to the phase just prior to the receipt of the initiator detected error message and retry the operation.

6.5.4 FORMAT UNIT COMMAND: OPCODE = 04(hex)

6.5.4.1 FORMAT UNIT COMMAND STRUCTURE:

**Table 6-10
FORMAT UNIT COMMAND**

Byte	Bit 7	6	5	4	3	2	1	0		
0	OPCODE = 04 _H									
1	LUN = 0		FMTDAT		CMPLST		DEFECT LIST FORMAT			
2	DATA PATTERN									
3-4	(MSB)		INTERLEAVE = XX				(LSB)			
5	VU = 0		RESERVED = 0			F	L			

The FORMAT UNIT command (Table 6-10) assigns logical blocks to physical sectors for optimum sequential access (within the limitations of available spare blocks), skipping known defective areas; and optionally writes a repetitive (one byte) data pattern in every block. The Q200 is already formatted when shipped from the factory, with data being undefined.

Issuing this command typically causes loss of data, even when the repetitive data pattern option is not used. The user should back up data prior to formatting.

FORMAT DATA (FMTDAT): A format data bit set to one indicates that defective block data is supplied during the DATA OUT phase. The defect list included with this data specifies the defects that shall be entered onto the defect map. The defect list refers to the current block length and current logical block addresses (not physical addresses). The format of the defect list is shown in Table 6-12. A FMTDAT bit set to zero indicates that the DATA OUT phase shall not occur (no defect list header and no defect data shall be supplied by the initiator). See Table 6-12.

COMPLETE LIST (CMPLST): A complete list bit set to one indicates that the data supplied is the complete list of known initiator-specified or field replacement defects. Any previous initiator-specified defect map or defect data shall be erased. The result is to purge any previous initiator-specified defect list and to build a new defect list which includes the initiator-specified list and the factory defects. A CMPLST bit of zero indicates that the data supplied is in addition to existing defect data using the current format.

DEFECT LIST FORMAT: The defect list format field specifies additional information related to the defect list. (See Table 6-11 for further information.)

DATA PATTERN: User data pattern to be written repetitively into each sector during FORMAT UNIT command (if option is enabled by the Fill Data Pattern bit in Page 39_H of the MODE SELECT command).

INTERLEAVE FIELD: Any value specified in this field by the initiator will be ignored by the Q200.

Table 6-11
FORMAT UNIT COMMAND VARIATIONS

Bit Reference (CDB byte 1)					Defect List Byte 1	Defect List Length	Option (See 6.5.4.2)
4	3	2	1	0	7	6	
1	1	0	X	X	1	1	0
1	1	0	X	X	0	0	0
1	0	0	X	X	1	1	0
0	X	X	X	X	NO DEFECT LIST*		
1	1	0	X	X	1	1	5
1	1	0	X	X	0	0	6
1	0	0	X	X	1	1	7
1	0	0	X	X	0	0	8
					<u>COMMENTS</u>		
1	X	1	0	X	Not Supported		
1	X	1	1	0	Vendor Unique (Not Used)		
1	X	1	1	1	Reserved		

Key: X = 1 or 0 (ignored by Q200)

*Because the DATA OUT phase does not occur, there is no defect list header.

The defect list shown in Table 6-12 contains a four-byte header followed by one or more defect descriptors. The defect list length specifies the total length in bytes of the defect descriptors that follow. Each defect descriptor for the logical block format specifies a four-byte defect logical block address that contains the defect. The defect list length is, therefore, equal to four times the number of defect descriptors. The defect addresses must be in ascending order. The defect list length may be zero.

Table 6-12
DEFECT LIST - LOGICAL BLOCK FORMAT

Byte	Bit	7	6	5	4	3	2	1	0
DEFECT LIST HEADER:									
0									
RESERVED = 0									
1									
FOV DPRY DCRT STPF RESERVED = 0									
2-3 (MSB) DEFECT LIST LENGTH (LSB)									
Defect Descriptor(s):									
0-4 (MSB) DEFECT LOGICAL BLOCK ADDRESS (LSB)									

FORMAT OPTIONS VALID (FOV): A format options valid bit set to one authorizes the setting of bits 4 through 6 (STPF, DCRT, and DPRY). An FOV bit set to zero indicates that the initiator requests the Q200 Series™ default format scheme to be implemented for the functions defined by bits 4 through 6.

DISABLE PRIMARY (DPRY): A disable primary bit set to one and an FOV bit set to one indicate that the Q200 shall exclude the factory defect map (primary list or P-list) from its list of flaws to manage while formatting. A DPRY bit set to zero indicates that the Q200 shall include the factory defect map in its list of flaws to manage while formatting. Note that the inclusion of the factory defect map is the default condition of the Q200.

DISABLE CERTIFICATION (DCRT): Not supported by the Q200. Must be set to zero by the initiator.

STOP FORMAT (STPF): Not supported by the Q200. Must be set to zero by the initiator.

6.5.4.2 APPLICATION OF FORMAT UNIT COMMAND

The Q200 offers eight formatting options as follows:

1. Format with absolutely no defects.
2. Format with original factory defects only.
3. Format with grown defects only (discarding factory defects).
4. Format with existing defects (factory and grown defects).
5. Format with provided defects only (discarding factory and existing grown defects).
6. Format with provided defects in addition to factory defects (discarding existing grown defects).
7. Format with provided defects in addition to existing grown defects (discarding factory defects).
8. Format with provided defects in addition to existing defects.

It should be noted that in any of the above options, the list of grown defects is always updated to reflect the latest condition; meanwhile, the original factory defects are always preserved.

The following tables illustrate the setting of the FORMAT UNIT CDB and defect list header required to implement each of the listed options.

(1) Format with absolutely no defects.

FORMAT UNIT Command

Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04H							
1	LUN = 0			1	1	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE = XX						(LSB)
5	VU = 0	RESERVED = 0				F	L	

Defect List Header

Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	1	1	0	0	RESERVED = 0			
2-3	(MSB)	DEFECT LENGTH = 0						(LSB)

NOTE: The list of grown defects will be erased.

(2) Format with original factory defects only.

FORMAT UNIT Command

Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04H							
1	LUN = 0			1	1	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE = XX						(LSB)
5	VU = 0	RESERVED = 0				F	L	

Defect List Header

Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	0	0	0	0	RESERVED = 0			
2-3	(MSB)	DEFECT LENGTH = 0						(LSB)

NOTE: The list of grown defects will be erased.

(3) Format with grown defects only; discard factory defects.

FORMAT UNIT Command

Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04 _H							
1	LUN = 0		1	0	0			
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE = XX						(LSB)
5	VU = 0	RESERVED = 0				F	L	

Defect List Header

Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	1	1	0	0	RESERVED = 0			
2-3	(MSB)	DEFECT LENGTH = 0						(LSB)

NOTE: The list of grown defects will be erased.

(4) Format with existing defects (factory and grown defects). The data out phase is not used.

FORMAT UNIT Command

Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04 _H							
1	LUN = 0		0	0	0			
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE = XX						(LSB)
5	VU = 0	RESERVED = 0				F	L	

NOTE: No defect list header or defect descriptor blocks follow the command.

(5) Format with provided defects only; discard factory and existing grown defects.

FORMAT UNIT Command

Bit Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04 _H							
1	LUN = 0			1	1	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE = XX					(LSB)	
5	VU = 0	RESERVED = 0				F	L	

Defect List Header

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	1	1	0	0	RESERVED = 0			
2-3	(MSB)	DEFECT LENGTH					(LSB)	

NOTE: The defect descriptor blocks follow the defect list header.

(6) Format with provided defects in addition to factory defects; discard existing grown defects)

FORMAT UNIT Command

Bit Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04 _H							
1	LUN = 0			1	1	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE = XX					(LSB)	
5	VU = 0	RESERVED = 0				F	L	

Defect List Header

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	0	0	0	0	RESERVED = 0			
2-3	(MSB)	DEFECT LENGTH					(LSB)	

NOTE: The defect descriptor blocks follow the defect list header. The provided defects will replace the existing grown defects.

(7) Format with provided defects in addition to existing grown defects; discard factory defects)

FORMAT UNIT Command

Bit Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04H							
1	LUN = 0			1	0	0		
2	DATA PATTERN							
3-4	(MSB)		INTERLEAVE = XX				(LSB)	
5	VU = 0		RESERVED = 0				F	L

Defect List Header

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	1	1	0	0	RESERVED = 0			
2-3	(MSB)		DEFECT LENGTH = 0					(LSB)

NOTE: The defect descriptor blocks follow the defect list header. The existing list of grown defects will be added with the provided defects, becoming the updated grown defects.

(8) Format with provided defects in addition to existing defects

FORMAT UNIT Command

Bit Byte	7	6	5	4	3	2	1	0
0	OPCODE = 04H							
1	LUN = 0			1	0	0		
2	DATA PATTERN							
3-4	(MSB)		INTERLEAVE = XX				(LSB)	
5	VU = 0		RESERVED = 0				F	L

Defect List Header

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	0	0	0	0	RESERVED = 0			
2-3	(MSB)		DEFECT LENGTH = 0					(LSB)

NOTE: The defect descriptor blocks follow the defect list header. The existing list of grown defects will be added with the provided defects, becoming the updated grown defects.

6.5.5 REASSIGN BLOCKS COMMAND: OPCODE = 07(hex)

Table 6-13
REASSIGN BLOCKS COMMAND

Bit Byte	7	6	5	4	3	2	1	0
0	OPCODE = 07 _H							
1	LUN = 0			RESERVED = 0				
2-4	RESERVED = 0							
5	VU = 0	RESERVED = 0			F	L		

The REASSIGN BLOCKS command (Table 6-13) requests the Q200 drive to reassign the defective logical blocks to areas on the drive reserved for this purpose.

The initiator transfers a defect list, during the DATA OUT phase following the command bytes, that contains the logical block addresses to be reassigned by the Q200. The drive will reassign the physical medium used for each logical block address in the list. The data contained in the logical blocks specified in the defect list will be preserved if recoverable within the limits prescribed by MODE SELECT error recovery parameters, and the data in all other logical blocks on the medium will be preserved. The effect of specifying a logical block to be reassigned that previously has been reassigned is to reassign the block again. The defective block address will be added to the drive's field replacement defect map.

The REASSIGN BLOCKS command should not be used to redundantly reassign defective logical blocks which have already been reassigned with the automatic reallocation feature of MODE SELECT.

Block reassignment can degrade performance by requiring access of spare locations. This can be remedied by moving a file to a more appropriate area (as determined using the READ CAPACITY command) or possibly by reassigning all logical block addresses with the FORMAT UNIT command.

The REASSIGN BLOCKS defect list (Table 6-14) contains a four-byte header followed by one or more defect descriptors. The length of each defect descriptor is four bytes.

The defect list length field specifies the total length in bytes of the defect descriptors that follow. The defect list length is, therefore, equal to four times the number of defect descriptors.

Table 6-14
REASSIGN BLOCKS DEFECT LIST

Byte	7	6	5	4	3	2	1	0		
Defect List Header:										
0-1	RESERVED = 0									
2-3	(MSB) DEFECT LIST LENGTH						(LSB)			
Defect Descriptor(s):										
0-4	(MSB) DEFECT LOGICAL BLOCK ADDRESS									

The defect descriptor specifies a four-byte defect logical block address that contains the defect. The defect addresses must be in ascending order. All of the 512-byte sectors associated with a logical block (1, 2, or 4 sectors) will be reassigned.

If the drive has insufficient capacity to reassign all of the defective logical blocks, the command will terminate with a CHECK CONDITION status and the sense key shall be set to MEDIUM ERROR. The logical block address of the first logical block not reassigned shall be returned in the information bytes of the sense data.

6.5.6 READ COMMAND: OPCODE = 08(hex)

6.5.6.1 READ COMMAND STRUCTURE:

Table 6-15
READ COMMAND

Byte	7	6	5	4	3	2	1	0						
0	OPCODE = 08H													
1	LUN = 0		(MSB) LOGICAL BLOCK ADDRESS											
2-3	LOGICAL BLOCK ADDRESS													
4	TRANSFER LENGTH													
5	VU = 0		RESERVED = 0				F	L						

The READ command (Table 6-15) requests that the target transfer data to the initiator.

LOGICAL BLOCK ADDRESS: The logical block address specifies the logical block at which the READ operation will begin.

TRANSFER LENGTH: The transfer length specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero indicates that 256 blocks shall be transferred. Acceptable values are 00-FFH.

6.5.6.2 READ COMMAND OPERATION

Read operations on the Q200 Intelligent Disk Drives are buffered. Each standard Q200 Series drive contains a full track (16K byte) buffer. The Q200 Series with DisCacheTM contains a 64K byte buffer. Because some of the buffer is reserved for internal use, the effective buffer size is slightly smaller (14K bytes for standard Q200 Series drive and 60K bytes for Q200s with DisCacheTM). When data is read from the drive, it is first stored in the buffer. Only after a logical block has been read successfully into the buffer are those data bytes transferred from the buffer out over the SCSI bus.

The two operations involved in reading, transferring data from the disk into the buffer and transferring data from the buffer to the SCSI bus, are essentially independent processes. The first operation, transferring data from the disk into the buffer, will continue as long as the transfer length indicates that there is more data to be read from the disk and there is room in the buffer. The second independent operation, transferring data from the buffer to the SCSI bus, will occur whenever the buffer contains at least one full logical block. If the buffer is completely filled because the second operation (buffer to SCSI bus) is slower than the first (disk to buffer), the first operation will be temporarily halted to allow the second operation to free-up room in the buffer. The first operation (disk to buffer) continues until the transfer length is exhausted. The second operation (buffer to SCSI bus) continues until the first operation is completed and the buffer is totally emptied.

The disconnect/reconnect strategy is tied to the dual processes of filling and emptying the buffer. The Q200 will not disconnect unless the initiator sends an IDENTIFY message with the "I Support Disconnect" bit (bit 6) set to one. If the disconnect is supported by the initiator, the drive will disconnect whenever the buffer is empty and the transfer count is not exhausted. The point at which reconnect occurs is controlled by the Buffer Full Ratio (MODE SELECT Page 2 byte 2). When the buffer fills to the level specified by the Buffer Full Ratio, the Q200 will initiate a reconnection. When reconnection is established, the buffer to SCSI bus transfer operation will resume. If there is a delay before reconnection (e.g., the bus is busy), the disk to buffer operation will continue filling the buffer until it is full, or the number of blocks requested has been reached.

6.5.7 WRITE COMMAND: OPCODE = 0A(hex)

6.5.7.1 WRITE COMMAND STRUCTURE:

Table 6-16
WRITE COMMAND

Bit Byte	7	6	5	4	3	2	1	0						
0	OPCODE = 0A _H													
1	LUN = 0		(MSB)		LOGICAL BLOCK ADDRESS									
2-3	LOGICAL BLOCK ADDRESS (LSB)													
4	TRANSFER LENGTH													
5	VU = 0	RESERVED = 0			F	L								

The WRITE command (Table 6-16) requests that the data transferred to the target be written in the area specified by the logical block address.

LOGICAL BLOCK ADDRESS: The logical block address specifies the logical block at which the WRITE operation will begin.

TRANSFER LENGTH: The transfer length specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero indicates that 256 blocks shall be transferred. Acceptable values are 00-FFH.

6.5.7.2 WRITE COMMAND OPERATION

Write commands on the Q200 Series Intelligent Disk Drives are buffered. Each Q200 contains a full track buffer (16K bytes) buffer. The Q200 Series drives with DisCache contains a 64K-byte buffer. Because some of the buffer is reserved for internal use, the effective buffer size is slightly smaller (14K bytes for the standard Q200 and 60K bytes for Q200s with DisCache). To store data on the disk, it is first stored in the buffer. When the head is correctly positioned on the track and is over the appropriate data area, data will be written from the buffer to the disk.

The two operations involved in writing, transferring data from the SCSI bus into the buffer and transferring data from the buffer onto the disk, are essentially independent processes. The first operation, transferring data from the SCSI bus into the buffer, will continue as long as the transfer length indicates that there is more data to be transferred across the SCSI bus from the initiator and there is room in the buffer. If the buffer fills up completely, the transferring of data across the SCSI bus will temporarily halt. The second independent operation, transferring data from the buffer onto the disk, will occur as long as the buffer is not empty and the head is correctly positioned over the appropriate data area. The only interaction between these two operations involves the buffer level. The buffer to disk operation will stop if the SCSI bus to buffer transfer process cannot keep up and the buffer empties. In turn, the SCSI to buffer transfer operation will temporarily halt if the buffer to disk transfer process falls behind and the buffer fills completely and runs out of room.

The disconnect/reconnect strategy is tied to the dual processes of filling and emptying the buffer. The Q200 will not disconnect unless the initiator sends an IDENTIFY message with the "I Support Disconnect" bit (bit 6) set to one. If disconnect is supported by the initiator, the drive will disconnect whenever it can no longer transfer data from the SCSI bus into the buffer because the buffer is full (or if the transfer count is exhausted) and the buffer to disk transfer process is waiting for head positioning to be completed. The positioning requirement restricts disconnections to situations where a substantial delay is required to prevent the possibility of constant disconnections/reconnections when the SCSI bus to buffer transfer rate is faster than the buffer to disk transfer rate. The point at which reconnection occurs is controlled by the Buffer Empty Ratio (MODE SELECT Page 2 byte 3). This ratio determines how low the buffer level must get via data transfer to disk before reconnection is attempted. After a successful reconnection, the SCSI bus to buffer transfer operation will resume.

6.5.8 SEEK COMMAND: OPCODE = 0B(hex)

Table 6-17
SEEK COMMAND

Bit Byte	7	6	5	4	3	2	1	0						
0	OPCODE = 0BH													
1	LUN = 0		(MSB)		LOGICAL BLOCK ADDRESS									
2-3	LOGICAL BLOCK ADDRESS (LSB)													
4	RESERVED = 0													
5	VU = 0	RESERVED = 0				F	L							

The SEEK command (Table 6-17) requests that the drive seek to the specified logical block address.

6.5.9 INQUIRY COMMAND: OPCODE = 12(hex)

6.5.9.1 INQUIRY COMMAND STRUCTURE:

Table 6-18
INQUIRY COMMAND

Bit Byte	7	6	5	4	3	2	1	0					
0	OPCODE = 12 _H												
1	LUN = 0			RESERVED = 0									
2-3	RESERVED = 0												
4	ALLOCATION LENGTH (BYTES)												
5	VU = 0	RESERVED = 0				F	L						

The INQUIRY command (Table 6-18) requests that information regarding identification of the target be sent to the initiator.

ALLOCATION LENGTH: The allocation length specifies the maximum number of bytes that the initiator has allocated for returned INQUIRY data (Table 6-19). The value of 00 indicates that no data will be returned. The acceptable values are 00-FF_H. The target will terminate the DATA IN phase when the allocation length bytes have been transferred or when all available INQUIRY data have been transferred to the initiator, whichever is less. (For the Q200, an allocation length of 56 (38_H) is recommended.)

6.5.9.2 INQUIRY DATA FORMAT:

Table 6-19
INQUIRY DATA FORMAT

Bit Byte	7	6	5	4	3	2	1	0							
0	PERIPHERAL DEVICE TYPE = 0														
1	RMB = 0	DEVICE TYPE QUALIFIER = 0													
2	RESERVED = 0	ECMA VERSION = 0			ANSI VERSION = 1										
3	RESERVED = 0			RESPONSE DATA FORMAT = 1											
4	ADDITIONAL LIST LENGTH														

BYTE	Q200 DRIVE IDENTIFICATION INFORMATION
5	VENDOR UNIQUE = 0
6 - 7	RESERVED = 0
8 - 15	VENDOR IDENTIFICATION = 'QUANTUM'
16 - 21	PRODUCT IDENTIFICATION (MODEL) = 'Q280' Q280 will be replaced by the correct model type
22 - 31	PRODUCT IDENTIFICATION (PART NUMBER) = '76-XXXXXX' XXXXXX is replaced by the appropriate digits
32 - 35	DRIVE REVISION LEVEL (VARIATION CODE) = 'VV' VV is replaced by the appropriate variation code
36 - 43	MICROCODE REVISION LEVEL = 'MMDDYYRR' MM = month DD = day YY = year RR = internal control code
44 - 55	DRIVE SERIAL NUMBER = 'yyddd-ssssM' yy = year drive built ddd = day of year drive built (Julian date) ssss = drive sequence number
	<u>NOTE:</u> The serial number may contain an "M" suffix indicating where the device was built.
56 - 97	RESERVED = 0
98 - N	VENDOR UNIQUE Reserved for factory use (Maximum N = 224)

If the drive is not ready when the INQUIRY command is received, it will be unable to return the correct data for bytes 22-31, 32-35 and 44-55. In this case, the data returned will be the default data as shown in the following table. Bytes 8 through 55 are ASCII values.

BYTES	DATA
8-15	"QUANTUM"
16-21	"Q280 "
22-31	"PART NUM "
32-35	"VCOD"
36-43	"CODE REV"
44-55	"DRV SER NUM "

where "Q280 " is replaced by "Q250 " for a Q250 drive.

PERIPHERAL DEVICE TYPE: The peripheral device type is set to zero for the Q200 disk drive (Direct-Access Device), unless the LUN specified in the Command Descriptor Block is invalid. In this case, the peripheral device type is set to $7F_H$. The rest of the data bytes are not affected and are returned as if the LUN was set to zero.

REMOVABLE MEDIUM BIT (RMB): The removable medium bit is always set to zero for the Q200 to indicate non-removable media.

DEVICE TYPE QUALIFIER: The device type qualifier is not supported by the Q200; it is set to zero.

VERSION: The ECMA version field is set to zero to indicate that the Q200 does not claim compliance to the European Computer Manufacturers Association version of SCSI.

The ANSI version field is set to 1_H to indicate compliance to the ANSI specification for SCSI.

RESPONSE DATA FORMAT: The response data format field is set to 1_H to indicate compliance to the Common Command Set proposed by an ANSC X3T9.2 working group.

ADDITIONAL LIST LENGTH: The additional list length specifies the length in bytes of the following list of drive identification information. If the allocation length of the command descriptor block is too small to transfer all of the vendor unique parameters, the additional length value will not be adjusted to reflect the truncation.

6.5.10 MODE SELECT COMMAND: OPCODE = 15(hex)

6.5.10.1 MODE SELECT COMMAND STRUCTURE:

**Table 6-20
MODE SELECT COMMAND**

Bit Byte	7	6	5	4	3	2	1	0						
0	OPCODE = 15_H													
1	LUN = 0		RESERVED = 0				SP							
2	RESERVED = 0													
3	RESERVED = 0													
4	PARAMETER LIST LENGTH (BYTES)													
5	VU = 0	RESERVED = 0				F		L						

The MODE SELECT command (Table 6-20) provides a means for the initiator to specify device parameters to the drive.

A MODE SELECT command will override any previous selection of device parameters, even from another initiator. Execution of the command will create a UNIT ATTENTION condition for other initiators.

When the Q200 drive is powered on and the motor spins up, the most recent set of device parameters is read from an unaddressable area on the drive (Saved Mode Table, see Section 6.4) and a UNIT ATTENTION condition is set for all

initiators. If the parameters cannot be successfully read from the drive, they will revert to their default values as shipped from the factory and the additional sense code will be set to MODE SELECT parameters have been changed (2AH) as if another initiator altered the parameters. If the parameters can be successfully read, the additional sense code will be set to reset occurred (29H) to indicate a reset (except if bit 1, byte 2 of MODE SELECT, Page 39H is set, in which case the UNIT ATTENTION condition will not occur; see Section 6.5.10.5).

PAGE FORMAT (PF) BIT: The page format bit set to one indicates that the data sent by the initiator after the MODE SELECT header and the block descriptors (if any) complies with the page format. If this bit is set to zero, then the data is vendor unique.

SAVE PARAMETERS (SP) BIT: If the save parameters bit is set to one, only the saveable pages of the Current Mode Table will be copied to the Saved Mode Table (see Section 6.4).

The parameter list length specifies the length in bytes of the MODE SELECT parameter list that shall be transferred during the DATA OUT phase. A parameter list length of zero indicates that no data shall be transferred. This condition shall not be considered as an error.

The MODE SELECT parameter list (Table 6-21) contains a four-byte header, followed by zero or one block descriptor, followed by zero or more "pages".

6.5.10.2 MODE SELECT PARAMETER LIST

Table 6-21
MODE SELECT PARAMETER LIST

Byte	7	6	5	4	3	2	1	0
Mode Select Header:								
0								RESERVED = 0
1								MEDIUM TYPE = 0
2								RESERVED = 0
3								BLOCK DESCRIPTOR LENGTH (BYTES)
Block Descriptors								
0								DENSITY CODE = 0
1-3	(MSB)							NUMBER OF BLOCKS = 0 (LSB)
4								RESERVED = 0
5-7	(MSB)							BLOCK LENGTH (LSB)
Page Descriptor(s):								
0	RESERVED = 0							PAGE CODE
1								PAGE LENGTH
2 to n								REFER TO PAGE DESCRIPTION

MEDIUM TYPE: The medium type is always set to zero which stands for default (currently mounted medium type) since the Q200 has non-removable media.

BLOCK DESCRIPTOR LENGTH: The block descriptor length specifies the length in bytes of all the block descriptors. It is equal to the number of block descriptors times eight and does not include the pages, if any. A block descriptor length of zero indicates that no block descriptors shall be included in the parameter list. This condition shall not be considered as an error. Only a single block descriptor is used for the Q200.

The block descriptor specifies the logical block length for the drive. The block descriptor contains a density code, a number of blocks, and a block length.

DENSITY CODE: The density code is always set to zero which stands for the default density of the medium, since the Q200 uses only one density.

NUMBER OF BLOCKS: The number of blocks field specifies the number of logical blocks on the medium that meet the density code and block length in the block descriptor. A number of blocks of zero indicates that all of the remaining logical blocks of the drive have the medium characteristics specified by the block descriptor. For the Q200, this parameter is always zero.

BLOCK LENGTH: The block length specifies the length in bytes of each logical block described by the block descriptor. For the Q200 allowable values are 512 (200_H), 1024 (400_H), or 2048 (800_H). When shipped from the factory the block length is set to 512.

Note that the block length may be altered without reformatting the drive or affecting stored data; however, logical block addresses corresponding to the data will be altered.

PAGE CODE: Pages may optionally be included (immediately following the block descriptor, if any) in any order. To avoid having to specify all Mode Parameters each time a MODE SELECT command is issued, the Mode Parameters are divided up into pages. A page is the minimum unit that can be specified in a MODE SELECT or MODE SENSE command. Each time a page is referenced, all of the parameters of that page must be specified. Each page contains parameters which are grouped together by common functionality. For example, Page 1 contains the Read/Write Error Recovery Parameters. These are the group of parameters which determine drive behavior during a data handling error. The Retry Count Parameter is located in this page, as are the bits which turn on/off error detection, determine if soft errors are reported, etc. Each page is numbered for reference. The Q200 supports the following pages:

<u>Page</u>	<u>Description</u>
01 _H	Read/Write Error Recovery Parameters
02 _H	Disconnect/Reconnect Control Parameters
03 _H	Direct Access Device Format Parameters (only via MODE SENSE)
04 _H	Rigid Disk Drive Geometry Parameters (only via MODE SENSE)
38 _H	Quantum Unique Cache Control Parameters
39 _H	Quantum Unique Drive Control Parameters

Pages 03_H and 04_H are read only pages available with the MODE SENSE command. The specification of these page codes in a MODE SELECT command will result in a CHECK CONDITION status with an ILLEGAL REQUEST sense key.

PAGE LENGTH: The page length indicates the number of bytes of the page that will follow, starting with the first byte of flags or values and continuing consecutively. The page length shall be set to the exact same value returned by the Q200 in the MODE SENSE Page Length byte. Otherwise, a CHECK CONDITION status with a sense key of ILLEGAL REQUEST will result.

6.5.10.3 ERROR RECOVERY PARAMETERS, Page Code 1H:

Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0				PAGE CODE = 1H			
1				PAGE LENGTH = 6H				
2	AWRE	ARRE	TB	RC	EEC	PER	DTE	DCR
3				RETRY COUNT				
4				CORRECTION SPAN				
5				HEAD OFFSET COUNT = 0				
6				DATA STROBE OFFSET COUNT = 0				
7				RECOVERY TIME LIMIT = 0				

DCR: A DCR (Disable Correction) bit 0 set to the value of one indicates that the data shall be transferred without correction, whether or not it is actually possible to correct the data. A DCR bit set to zero indicates that the data shall be corrected if possible. Uncorrectable data is transferred with no attempt to apply correction. If RC is set to one, the DCR bit is ignored. The default value is zero.

DTE: A DTE (Disable Transfer on Error) bit 1 set to the value of one indicates that the Q200 shall create the CHECK CONDITION status immediately upon detection of an error, if PER is set to one. The data transfer to the initiator shall be terminated. The block in error may or may not be transferred to the initiator depending upon the setting of the TB bit. A DTE bit set to zero enables data transfer for any data which can be recovered within the limits of the Error Recovery Flags. Any errors that would be posted are not posted until the transfer length is exhausted. If PER is set to zero, or RC is set to one, the DTE bit is ignored. The default value is zero.

PER: A PER (Post Error) bit 2 set to the value of zero indicates that the Q200 shall not create the CHECK CONDITION status for errors recovered within the limits established by the other Error Recovery Flags. Recovery procedures exceeding the limits established by the other Error Recovery Flags shall be posted accordingly by the Q200. The transfer may terminate prior to exhausting the transfer length depending on the error and state of the other Error Recovery Flags. A PER bit set to one shall enable the reporting of the CHECK CONDITION status for detected errors, with the appropriate sense key. If multiple errors occur, the sense data shall report the block address of either the unrecoverable error, or if no unrecoverable error occurred, the block with the last recovered error. If the RC bit is set to one, the PER bit is ignored. The default value is zero.

EEC: An EEC (Enable Early Correction) bit 3 set to the value of one indicates that the Q200 will use its ECC algorithm immediately upon detection of a data error (i.e., a non-zero syndrome). The drive will not perform read retries before applying correction unless the error is determined to be uncorrectable. Seek or positioning retries and the recovery procedure retries of the message system are not affected by the value of this bit. An EEC bit set to zero indicates that the Q200 will use its normal recovery procedures when an error occurs (see Appendix A). If the RC bit is one, the ECC bit is ignored. The default value is zero.

RC: An RC (Read Continuous) bit 4 set to the value of one requests the Q200 to transfer the requested length without adding delays which would increase or ensure data integrity (i.e., delays caused by the Q200 error recovery schemes). This implies that the Q200 shall send data which may be erroneous in order to maintain a continuous flow of data and avoid delays, and that ignored errors will create the CHECK CONDITION status. An RC bit set to zero indicates that time consuming operations for error recovery are acceptable during the data transfer. The default value is zero.

TB: A TB (Transfer Block) bit 5 set to the value of one, when the DTE bit is set to one, indicates that the failing block data (recovered or unrecoverable) shall be transferred to the initiator. A TB bit set to zero, when DTE bit is set to one, indicates that the failing block data shall not be transferred to the initiator. If DTE is zero, or RC is one, the TB bit is ignored. The default value is zero.

AARE: An ARRE (Automatic Read Reallocation Enabled) bit 6 set to the value of one indicates that the target shall enable automatic reallocation of bad blocks. The execution of the automatic reallocation is similar to the function of the REASSIGN BLOCKS command, but is initiated by the Q200 when a hard error is encountered (i.e., same non-zero ECC syndrome on two consecutive retries. See Appendix A). An ARRE bit set to zero indicates that the target shall not perform automatic reallocation of bad blocks. If RC is set to one, the ARRE bit is ignored. The default value is zero.

The AWRE (Automatic Write Reallocation Enabled) bit is not supported by the Q200. It must be set to zero by the initiator.

Retry count is the number of times that the Q200 shall attempt to recover from data errors by rereading before applying correction. The default value is 8. This parameter is ignored if both the DTE and PER bits are set one, if the RC bit is set to one, or if the EEC bit is set to one.

Correction span is not implemented on the Q200.

Data strobe offset is not implemented on the Q200.

Recovery time limit is not implemented on the Q200.

The following table summarizes the valid modes of operation on the Q200.

EEC	PER	DTE	DCR	Description
0	0	0	0	Normal error recovery procedure. Read retries are attempted until good data is read, a stable syndrome is obtained, or the retry count is exhausted. If needed, ECC is invoked if correction is possible. Complete data transfer occurs unless an uncorrectable error is encountered. Only uncorrectable errors are reported.
0	0	0	1	Same as (0,0,0,0) except no ECC correction is attempted. If read retries are not successful, data transfer is stopped and an unrecoverable error is reported.
0	0	1	0	Invalid Request.
0	0	1	1	Invalid Request.
0	1	0	0	Same as (0,0,0,0) except all data errors (recoverable and unrecoverable) are reported. A recoverable error is reported after the data transfer is completed.
0	1	0	1	Same as (0,0,0,0) except all data errors are reported. A data error recovered via read retries is reported after the data transfer is completed.
0	1	1	0	Read retries are attempted until good data is read, stable syndrome is obtained, or the retry count is exhausted. If needed, ECC is invoked if correction is possible. Data transfer is stopped upon error detection and all data errors are reported.
0	1	1	1	Same as (0,1,1,0) except no ECC correction is attempted. If read retries are not successful, the error is reported as unrecoverable.
1	0	0	0	ECC is invoked immediately if correction is possible. If error is uncorrectable, read retries are attempted until good data is read, stable syndrome is obtained, or the retry count is exhausted. Complete data transfer occurs unless an unrecoverable error is encountered. Only unrecoverable errors are reported.
1	0	0	1	Invalid Request.
1	0	1	0	Invalid Request.
1	0	1	1	Invalid Request.
1	1	0	0	Same as (1,0,0,0) except all data errors are reported. A recoverable error is reported after the data transfer is completed.
1	1	0	1	Invalid Request.
1	1	1	0	Same as (1,0,0,0) except data transfer is stopped upon error detection (recoverable and unrecoverable) and all data errors are reported.
1	1	1	1	Invalid Request.

6.5.10.4 DISCONNECT/RECONNECT CONTROL PARAMETERS, Page Code 2H:

Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							PAGE CODE = 2H
1								PAGE LENGTH = A _H
2								BUFER FULL RATIO
3								BUFER EMPTY RATIO
4-5	(MSB)							BUS INACTIVITY LIMIT = 0 (LSB)
6-7	(MSB)							DISCONNECT TIME LIMIT = 0 (LSB)
8-9	(MSB)							CONNECT TIME LIMIT = 0 (LSB)
10-11								RESERVED

BUFFER FULL RATIO: indicates to the Q200, for commands requiring data transfer to the initiator, how full the buffer shall be prior to reconnecting (unless the buffer contains all of the required data). The value represents the percentage of the buffer that will fill; i.e., 255 (FF_H) represents 100%, 128 (80_H) represents 50% and so on. Each bit is equal to 1/256th of the maximum buffer size. The default value is 0 (00_H) which is interpreted to mean that at least 512 bytes must be transferred to the buffer before reconnection to the initiator.

BUFFER EMPTY RATIO: indicates to the Q200, for commands requiring data transfer from the initiator, how empty the buffer shall be prior to reconnecting to fetch more data (unless the buffer can contain all of the required data). The value represents a quantity similar to the buffer full value; i.e., each bit represents 1/256th of the maximum buffer size. The value represents the percentage of the buffer that will empty; i.e., 255 (FF_H) represents 100% empty, 128 (80_H) represents 50% and so on. The default value is 0 (00_H) which is interpreted to mean that there is space for at least 512 bytes of data in the buffer before reconnection to the initiator.

Both the Buffer Full Ratio and the Buffer Empty Ratio are rounded (down or up, respectively) by the Q200 to the nearest multiple of 512 bytes for commands requiring logical block transfer. None of the remaining bytes (bytes 4 through 11) are implemented on the Q200.

6.5.10.5 QUANTUM UNIQUE CACHE CONTROL PARAMETERS, Page Code 38H:

Bit Byte	7	6	5	4	3	2	1	0				
0	PAGE NUMBER = 38H											
1	PAGE LENGTH = 0EH											
2	RESVD	WIE	RESVD	CE	CACHE TABLE SIZE							
3	PREFETCH THRESHOLD											
4	MAXIMUM PREFETCH - ABSOLUTE											
5	MAXIMUM PREFETCH - MULTIPLE											
6	MINIMUM PREFETCH - ABSOLUTE											
7	MINIMUM PREFETCH - MULTIPLE											
8-15	RESERVED											

WRITE INDEX ENABLE (WIE): All data from WRITES create a cache segment stored in the cache and indexed by the cache table. A WIE (Write Index Enable) bit 6 set to the value of zero indicates that the Q200 will overwrite this segment in the cache on the next READ or WRITE. Thus, this segment is immediately removed from the cache. A WIE bit set to one indicates that the Q200 will treat this segment as if it had been generated by a READ. Thus, the FIFO replacement scheme will be applied to this segment (as it is to all other segments). A WIE bit set to zero preserves more data in the cache from READs, often increasing the likelihood that the requested data lies in the cache (the magnitude of this increase depends on the size and frequency of WRITE requests). The default value is one.

A CE (Cache Enable) bit 4 set to the value of one indicates that the Q200 will have caching operational. A CE bit set to zero indicates that the Q200 will disable caching and use the 60K RAM as a transfer buffer only. The default value is 1 (cache enabled).

Cache Table Size bits 0 through 3 indicate to the Q200 the number of segments to index in the cache (i.e., the number of entries in the cache table). The valid range is 1 through 12; any value outside this range will result in a CHECK CONDITION status with a sense key of ILLEGAL REQUEST. The default value is 12.

Prefetch Threshold indicates to the Q200 the maximum transfer length, in logical blocks, of a READ command for which a prefetch operation will be performed. If the transfer length exceeds this value, no prefetch will occur. The valid range is 0 to 116; any value outside this range will result in a CHECK CONDITION status with a sense key of ILLEGAL REQUEST. The default value is 16.

Maximum Prefetch - Absolute - indicates to the Q200, when added to the product of Maximum Prefetch - Multiple and the READ transfer length, the maximum number of logical blocks to prefetch on a cache miss for a READ command. The valid range is 0 to 116, any value outside this range will result in a CHECK CONDITION status with a sense key of ILLEGAL REQUEST. The default value is 0.

Maximum Prefetch - Multiple - indicates to the Q200, when first multiplied by the READ transfer length and then added to the Maximum Prefetch - Absolute value, the maximum number of logical blocks to prefetch on a cache miss for a READ command. The valid range is 0 to 116; any value outside this range will result in a CHECK CONDITION status with a sense key of ILLEGAL REQUEST. The default value is 3. For example,

with a READ transfer length of 2K bytes, the Maximum Prefetch -Multiple default value of 3, and the Maximum Prefetch -Absolute default value 0, the maximum prefetch amount is (2K bytes X 3) + 0 = 6K bytes.

NOTE: The *DisCache* memory size is 60K bytes but the largest single segment is limited to 116, 512-byte sectors. If the calculated maximum prefetch size is greater than 116 512-byte sectors, the Q200 will only prefetch 116 sectors. This insures that every prefetch results in valid data in the cache.

Minimum Prefetch - Absolute indicates to the Q200, when added to the product of Minimum Prefetch - Multiple and READ transfer length, the minimum number of logical blocks to prefetch on a cache miss for a READ command. The valid range is 0 to 116; any value outside this range will result in a CHECK CONDITION status with a sense key of ILLEGAL REQUEST. The default is 0.

Minimum Prefetch - Multiple - indicates to the Q200, when first multiplied by the READ transfer length and then added to the Minimum Prefetch - Absolute value, the minimum number of logical blocks to prefetch on a cache miss for a READ command. The valid range is 0 to 116; any value outside this range will result in a CHECK CONDITION status with a sense key of ILLEGAL REQUEST. The default value is 0.

NOTE: The minimum prefetch default values of 0 insure that the prefetch operation will be terminated upon reading the current sector whenever the host issues another request to the Q200.

6.5.10.6 QUANTUM UNIQUE DRIVE CONTROL PARAMETERS, Page Code 39H:

Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0				PAGE CODE = 39H			
1				PAGE LENGTH = 6H				
2	DIO	DII	FDB	RUEE	FDPE	R=0	DUA	DRT
3	DDIS	DELDIS	RESERVED = 0	SSID		SCSIADR		
4-7			RESERVED = 0					

A DRT (Disable Reselection Timeout) bit 0 set to the value of one indicates that the Q200 will not timeout during a reselection request. A DRT bit set to zero indicates that the Q200 will clear the bus if there is no BSY response from the initiator after waiting a selection time out period (250 msec). The default value is zero.

A DUA (Disable Unit Attention) bit 1 set to the value of one indicates that the Q200 will not issue a CHECK CONDITION status with the UNIT ATTENTION sense key on the first command after a reset. The first command issued after the reset condition will be executed as usual. A DUA bit set to zero indicates that the unit attention condition will exist after the Q200 has been powered-on or reset (by a BUS DEVICE RESET message or a "hard" RESET condition). The default value is zero.

Bit 2 is reserved; it must be set to zero.

An FDPE (Fill Data Pattern Enabled) bit 3 set to the value of one indicates that the Q200 will write the data pattern specified in byte 2 of the FORMAT UNIT command into every user-accessible sector of the drive when a FORMAT UNIT command is executed. An FDPE bit set to zero indicates that the Q200 will ignore the information in byte 2 of the FORMAT UNIT command. Issuing a FORMAT UNIT command will typically cause loss of data, even if this bit is not enabled. The user should back up data prior to formatting. The default value is zero.

An **RUEE** (Reallocate Uncorrectable Error Enabled) bit 4 set to the value of one indicates that if the ARRE bit in Page 1H (the Error Recovery Parameters) is set to one, the Q200 will automatically reallocate uncorrectable hard errors. The uncorrectable (bad) data is written to a new block and a new ECC (one that is correct for the bad data) is appended. This permits the new block to be read without error; however, the data is still the original uncorrectable or bad data. An RUEE bit set to zero indicates that the Q200 will not attempt automatic reallocation if the hard error encountered is determined to be uncorrectable. In this case, a CHECK CONDITION with a sense key of MEDIUM ERROR - Uncorrectable data error will result. The default value is zero.

An **FDB** (Fast Deassertion of Busy) bit 5 set to the value of zero indicates that the COMMAND COMPLETE message will be sent from the Q200 to the initiator immediately upon completion of the command. This introduces the possibility that the code sequence, from the initiation of the MESSAGE IN phase to the BUS FREE phase, will be interrupted for approximately 300 μ s by the Q200's processing of servo information. An FDB bit set to one can delay interruption by synchronizing the sending of the COMMAND COMPLETE message with the last servo interrupt. The Q200 will deassert BUSY approximately 45 μ s after the initiator asserts the ACKNOWLEDGE line for the COMMAND COMPLETE message. Therefore, the ACKNOWLEDGE turn around time should be less than 105 μ s to ensure that the time, from the Q200 placing the COMMAND COMPLETE message onto the bus to the time when BUSY is deasserted by the Q200, is less than 150 μ s. The default value is zero.

Bits 6 and 7 are reserved; they must be set to zero.

An **SSID** (Software Selectable ID) bit 3 set to the value of zero indicates that the drive SCSI ID will be determined by the value indicated by the jumpers (on the drive's PCB) at all times. An SSID bit set to one indicates that the drive SCSI ID will be determined by the value of the SCSIADR bits (bits 0-2 of byte 3), and the value indicated by the jumpers will be ignored. When the SSID bit is set to one, the SCSI ID will not be altered until the next RESET (hard or warm) or power up of the drive. If, upon initial power up, the SSID bit is set (i.e., it was previously set and saved to the disk), then the SCSIADR bits will determine the SCSI ID; otherwise, the SCSI ID will be that determined by the jumpers. The default value of SSID is zero. The default value of SCSIADR is zero.

A **DELDIS** (Delay of Disconnection) bit 6 set to the value of one prevents a delay of more than 15 μ s between any two bytes transferred during DATA phase of READ, WRITE and COPY commands, except at block boundaries. Disconnection and the response to the ATTENTION signal are limited to block boundaries. A DELDIS bit set to zero causes no action to occur. The default value is zero.

A **DDIS** (Disable Disconnection) bit 7 set to the value of one causes the suppression of disconnection and reconnection during a data transfer (READ and WRITE). The initial disconnection due to the implied seek of a READ command is NOT prevented, but no other disconnects will occur. The WRITE command will also allow the initial disconnection due to the seek and prevent any further disconnects; however, the WRITE command starts the data transfer from the host prior to initiating the seek so that the data transfer is already in progress at the time of the seek. This means that the drive may disconnect once at the start of the WRITE command (after the data transfer has started), but will not disconnect again. A DDIS bit set to zero causes no action to occur. The default value is zero.

6.5.11 RESERVE COMMAND: OPCODE = 16(hex)

Table 6-22
RESERVE COMMAND

Byte	Bit 7	6	5	4	3	2	1	0
0	OPCODE = 16H							
1	LUN = 0		3RDPTY		3RD PARTY DEVICE ID			EXTENT
2	RESERVATION IDENTIFICATION = XX							
3-4	(MSB)		EXTENT LIST LENGTH = XX			(LSB)		
5	VU=0		RESERVED = 0			F	L	

The RESERVE command (Table 6-22) is used to reserve a logical unit for the use of the initiator or another specified SCSI device. The Q200 is considered a single, indivisible logical unit with a LUN of zero.

LOGICAL UNIT RESERVATION: The extent bit must be set to zero. This command requests that the drive be reserved for the exclusive use of the initiator until the reservation is superseded by another valid RESERVE command from the same initiator or is released by a RELEASE command from the same initiator, by a BUS DEVICE RESET message from any initiator, or by a "hard" RESET condition. A logical unit reservation will not be granted if the drive is reserved by another initiator. It is permissible for an initiator to reserve the drive when it is currently reserved for that initiator. The reservation identification and the extent list length will be ignored.

If the drive is reserved for another initiator, it will respond by either returning a RESERVATION CONFLICT status or, if disconnection is allowed, by queuing the RESERVE request and disconnecting.

If, after honoring the reservation, any other initiator then subsequently attempts to perform any command on the reserved drive (other than a RELEASE command, which will be ignored or a RESERVE command, which may be queued), then the command shall be rejected with RESERVATION CONFLICT status.

EXTENT RESERVATION (EXTENT) SET TO ONE: Not supported for the Q200.

If the extent bit is set to one, the RESERVE command will be rejected with CHECK CONDITION status with the ILLEGAL REQUEST sense key.

THIRD-PARTY RESERVATION (3RDPTY): The third-party reservation option for the RESERVE command allows an initiator to reserve a logical unit for another SCSI device. This option is intended for use in multiple-initiator systems that use the COPY command.

If the third-party reservation bit is set to zero, then the third-party reservation option is not requested. If the third-party bit is set to one, then the RESERVE command reserves the drive for the SCSI device specified in the third-party device ID field. The drive shall preserve the reservation until it is superseded by another valid RESERVE command from the same initiator or until it is released by the same initiator, by a BUS DEVICE RESET message from any initiator, or by a "hard" RESET condition. The drive shall ignore any attempt to release the reservation made by any other initiator, including the specified third party.

6.5.12 RELEASE COMMAND: OPCODE = 17(hex)

Table 6-23
RELEASE COMMAND

Byte	7	6	5	4	3	2	1	0					
0	OPCODE = 17H												
1	LUN = 0		3RDPTY		3RD PARTY DEVICE ID			EXTENT					
2	RESERVATION IDENTIFICATION												
3-4	RESERVED = 0												
5	VU=0		RESERVED = 0			F	L						

The RELEASE command (Table 6-23) is used to release a previously reserved logical unit. It is not an error for an initiator to attempt to release a reservation that is not currently active.

LOGICAL UNIT RELEASE: The extent bit must be set to zero. This will cause the drive to terminate an active reservation from the initiator.

EXTENT RELEASE (EXTENT) SET TO ONE: Not supported for the Q200.

If the extent bit is set to one, the RELEASE command will be rejected with CHECK CONDITION status with the ILLEGAL REQUEST sense key.

THIRD-PARTY RESERVATION (3RDPTY): The third-party release option for the RELEASE command allows an initiator to release a logical unit that was previously reserved using the third-party reservation option. This option is intended for use in multiple-initiator systems that use the COPY command.

If the third-party bit is set to zero, then the third-party release option is not requested. If the third-party bit is set to one, then the target shall release the specified logical unit, but only if the reservation was made using the third-party reservation option by the same initiator for the same SCSI device as specified in the third-party device ID field.

6.5.13 COPY COMMAND: OPCODE = 18(hex)

6.5.13.1 COPY COMMAND STRUCTURE:

Table 6-24
COPY COMMAND

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPCODE = 18H								
1		LUN = 0				RESERVED = 0				
2-4		(MSB)		PARAMETER LIST LENGTH				(LSB)		
5		VU = 0			RESERVED = 0					

The COPY command (Table 6-24) provides a means to copy data from one logical unit to another. The Q200 drives cannot copy to themselves, nor can the Q200 drives perform a third party copy (i.e., a COPY command issued to the Q200 drive will not cause two other SCSI devices to copy between themselves). The Q200 drive must be either the source or destination unit, but not both.

The parameter list length specifies the length in bytes of the parameters that shall be sent during the DATA OUT phase of the command. A parameter list length of zero indicates that no data shall be transferred. This condition shall not be considered as an error.

The COPY parameter list (Table 6-25) begins with a four-byte header that contains the COPY function code and priority. Following the header is one or more segment descriptors.

6.5.13.2 COPY PARAMETER LIST

Table 6-25
COPY PARAMETER LIST

Byte	Bit	7	6	5	4	3	2	1	0			
0		COPY FUNCTION						PRIORITY = 1H				
1		VENDOR UNIQUE = 00										
2-3		RESERVED = 0										

SEGMENT DESCRIPTOR(S):

0-n	SEGMENT DESCRIPTOR 0 (SEE SPECIFIC TABLE FOR LENGTH)
-----	------------------------------------------------------

The COPY function codes (Table 6-26) define a specific format for the segment descriptors.

The priority field of the COPY parameter list establishes the relative priority of this COPY command to other commands being executed by the same target. All other commands are assumed to have a priority of 1. Priority 0 is the highest priority with increasing values indicating lower priorities. For the Q200, all commands have a priority of 1.

The segment descriptor formats (Tables 6-27 and 6-28) are determined by the COPY function code. A maximum of 256 segment descriptors are permitted, and are identified by ascending numbers beginning with zero.

Table 6-26
COPY FUNCTIONS

<u>Peripheral Device Type</u>		<u>COPY Function Code</u>	<u>Segment Descriptor Table</u>	<u>Comment</u>
<u>Source</u>	<u>Destination</u>			
00 _H	01 _H	00 _H	Table 6-27	Direct-Access
00 _H	02 _H	00 _H	Table 6-27	to
00 _H	03 _H	00 _H	Table 6-27	Sequential-Access
01 _H	00 _H	01 _H	Table 6-27	Sequential-Access
03 _H	00 _H	01 _H	Table 6-27	to Direct-Access
00 _H	00 _H	02 _H	Table 6-28	
00 _H	04 _H	02 _H	Table 6-28	
04 _H	00 _H	02 _H	Table 6-28	Direct-Access
05 _H	00 _H	02 _H	Table 6-28	to Direct-Access
Peripheral device type:				
		00 _H	Direct-access device (e.g., Q200)	
		01 _H	Sequential-access device	
		02 _H	Printer device	
		03 _H	Processor device	
		04 _H	Write-once read-multiple device	
		05 _H	Read-only direct-access device	
COPY function code:				
		00 _H	Direct-access to sequential-access	
		01 _H	Sequential-access to direct-access	
		02 _H	Direct-access to direct-access	
		03 _H	Not Used	
		04 _H - 0F _H	Reserved	
		10 _H - 1F _H	Vendor unique	

6.5.13.3 COPY FUNCTION CODE 00H AND 01H

The format for the segment descriptors for COPY transfers between direct-access and sequential-access devices is specified in Table 6-27. This format is required for COPY function codes 00H or 01H. The segment descriptor may be repeated up to 256 times within the parameter list length specified in the command descriptor block.

Table 6-27
SEGMENT DESCRIPTOR FOR COPY FUNCTION CODES 00H AND 01H

Byte	Bit	7	6	5	4	3	2	1	0
0		SOURCE ADDRESS			RESERVED = 0		SOURCE LUN		
1		DESTINATION ADDRESS			RESERVED = 0		DESTINATION LUN		
2-3	(MSB)	SEQUENTIAL-ACCESS DEVICE BLOCK LENGTH					(LSB)		
4-7	(MSB)	DIRECT-ACCESS DEVICE NUMBER OF BLOCKS					(LSB)		
8-11	(MSB)	DIRECT-ACCESS DEVICE LOGICAL BLOCK ADDRESS					(LSB)		

Source address and destination address fields specify the SCSI devices and the source LUN and destination LUN fields specify the logical units to use for this segment of the COPY command. If an unsupported COPY operation is requested, the command shall be terminated with a CHECK CONDITION status with the ILLEGAL REQUEST sense key.

The sequential-access device block-length field specifies the block-length to be used on the sequential-access logical unit during this segment of the COPY command. If this block-length is known, by the SCSI device managing the COPY, to be not supported, the command shall be terminated with a CHECK CONDITION status with the ILLEGAL REQUEST sense key. If the block-length is found to be invalid while executing a read or write operation to the sequential-access device, the command shall be terminated with a CHECK CONDITION status with the COPY ABORTED sense key.

The direct-access device number of blocks field specifies the number of blocks in the current segment. A value of zero indicates that no blocks shall be transferred in this segment. The direct-access device logical block address specifies the starting logical block address on the logical unit for this segment.

6.5.13.4 COPY FUNCTION CODE 02H

The format for the segment descriptors for COPY transfers among direct-access devices is specified by Table 6-28. This format is required for COPY function code 02H. The segment descriptor may be repeated up to 256 times within the parameter list length specified in the command descriptor block.

Table 6-28
SEGMENT DESCRIPTOR FOR COPY FUNCTION CODE 02H

Byte	7	6	5	4	3	2	1	0
0	SOURCE ADDRESS			RESERVED = 0		SOURCE LUN		
1	DESTINATION ADDRESS			RESERVED = 0		DESTINATION LUN		
2-3				RESERVED = 0				
4-7	(MSB)	SOURCE NUMBER OF BLOCKS					(LSB)	
8-11	(MSB)	SOURCE LOGICAL BLOCK ADDRESS					(LSB)	
12-15	(MSB)	DESTINATION LOGICAL BLOCK ADDRESS					(LSB)	

The **source address** and **destination address** fields specify the SCSI devices; the **source LUN** and **destination LUN** specify the logical units to use for this segment of the COPY command. If an unsupported COPY operation is requested, the command shall be terminated with a CHECK CONDITION status with the **ILLEGAL REQUEST** sense key.

The **source number of blocks** field specifies the number of blocks to be transferred from the source device during command execution. The **source logical block address** field specifies the starting logical block address on the source device. The **destination logical block address** field specifies the starting logical block address on the destination device.

6.5.13.5 ERRORS DETECTED DURING THE EXECUTION OF A COPY COMMAND

Two classes of unusual conditions may occur during the execution of a COPY command.

6.5.13.5.1 ERRORS DETECTED BY THE MANAGING SCSI DEVICE

The first class consists of unusual conditions detected by the SCSI device that received the COPY command and is managing the execution of the command. These conditions include parity errors while transferring the COPY command and status byte, invalid parameters in the COPY command, invalid segment descriptors, and the inability of the SCSI device controlling the COPY functions to continue operating. In the event of such an unusual condition, the SCSI device managing the COPY shall:

- (1) Terminate the COPY command with a CHECK CONDITION status.
- (2) Return the SENSE DATA in the extended sense format. The valid bit shall be set to one. The segment number shall contain the number of the segment descriptor being processed at the time the unusual condition is detected. The sense key shall contain the sense key code describing the unusual condition. The information bytes shall contain the difference between the number of blocks field in the segment descriptor being processed at the time of the failure and the number of blocks successfully copied. This number is the residue of unprocessed blocks remaining for the segment descriptor.

6.5.13.5.2 ERRORS DETECTED BY A TARGET

The second class of errors consists of unusual conditions detected by the SCSI device transferring data at the request of the SCSI device managing the transfer. The SCSI device managing the COPY command detects unusual conditions by receiving a CHECK CONDITION status from one of the SCSI devices it is managing. It shall then recover the SENSE DATA associated with the unusual condition.

The SCSI device managing the COPY command may also be the source or destination SCSI device (or both). It shall distinguish between a failure of the management of the COPY and a failure of the data transfer being requested. It then shall create the appropriate SENSE DATA internally.

After recovering the SENSE DATA associated with the detected error, the SCSI device managing the COPY command shall:

- (1) Terminate the COPY command with a CHECK CONDITION status.
- (2) Return the SENSE DATA in the extended sense format. The valid bit shall be set to one. The segment number shall contain the number of the segment descriptor being processed at the time the unusual condition is detected. The sense key shall be set to COPY ABORTED. The information bytes shall contain the difference between the number of blocks field in the segment descriptor being processed at the time of the failure and the number of blocks successfully copied. This number is the residue of unprocessed blocks remaining for the segment descriptor. The additional sense length shall specify the number of additional sense bytes.

The first additional sense byte (byte 8) shall specify the byte number, relative to the first byte of SENSE DATA (which is considered byte number zero) of the beginning of the source logical unit's status byte and SENSE DATA. A zero value indicates that no status byte or SENSE DATA is being returned for the source logical unit. The first byte of the area pointed to by the first additional sense byte shall contain the status byte from the source logical unit. The subsequent bytes shall contain, unchanged, the SENSE DATA recovered from the source logical unit.

The second additional sense byte (byte 9) shall specify the byte number, relative to the first byte of SENSE DATA (which is considered byte number zero) of the beginning of the destination logical unit's status byte and SENSE DATA. A zero value indicates that no status byte or SENSE DATA is being returned for the destination logical unit. The first byte of the area pointed to by the second additional sense byte shall contain the status byte from the destination logical unit. The subsequent bytes shall contain, unchanged, the SENSE DATA recovered from the destination logical unit.

6.5.14 MODE SENSE COMMAND: OPCODE = 1A(hex)

6.5.14.1 MODE SENSE COMMAND STRUCTURE:

Table 6-29
MODE SENSE COMMAND

Byte	Bit	7	6	5	4	3	2	1	0						
0		OPCODE = 1AH													
1		LUN = 0			RESERVED = 0										
2		PCF		PAGE CODE											
3		RESERVED = 0													
4		ALLOCATION LENGTH													
5		VU=0		RESERVED = 0			F	L							

The MODE SENSE command (Table 6-29) provides a means for the target to report its peripheral device parameters to the initiator. It is a complementary command to the MODE SELECT command.

PAGE CONTROL FIELD (PCF): The PCF (Page Control Field) bits 6 and 7 defines the type of Page Parameter values to be returned. There are four possible options.

Report Current Values (PCF = 00)

The page defined by the page code is returned to the initiator with fields and bits set to current values. The current values are either:

- as set in the last successfully completed MODE SELECT command.
- or are identical to the saved values if no MODE SELECT commands were issued since the last power-up or reset.

Fields and bits not supported by the Q200 are set to zero.

Report Changeable Values (PCF = 01)

The page defined by the page code is returned to the initiator with fields and bits which are changeable by the initiator set to all ones. Fields and bits not allowed to be changed by the initiator are set to zero.

Report Default Values (PCF = 10)

The page defined by the page code is returned to the initiator with fields and bits set to the Q200 SeriesTM default values (i.e., as shipped from the factory). Fields and bits not supported by the Q200 are set to zero. If there is confusion over whether a returned value of zero indicates a default parameter or a nonsupported parameter, the initiator can examine the changeable values.

Report Saved Values (PCF = 11)

The page defined by the page code is returned to the initiator with fields and bits set to saved values. The saved values are either:

- as set in the last successfully completed MODE SELECT command with the SP (Save Parameters) bit 0, byte 1 of the CDB set to one.
- or are identical to the default values if no MODE SELECT commands with the SP bit set were ever issued (and completed successfully).

Fields and bits not supported by the Q200 are set to zero.

PAGE CODE: For all of the above options, the page length byte value returned indicates the number of bytes the Q200 supports within each page. This same value must be specified in the page length field (byte 1 of each page descriptor) when issuing the MODE SELECT command.

The page code specifies the particular page information to be returned to the initiator in the MODE SENSE data. The Q200 supports the following pages:

Page	Description
01 _H	Read/Write Error Recovery Parameters
02 _H	Disconnect/Reconnect Control Parameters
03 _H	Direct-Access Device Format Parameters (Read Only; non-changeable)
04 _H	Rigid Disk Drive Geometry Parameters (Read Only; non-changeable)
38 _H	Quantum Unique Cache Control Parameters
39 _H	Quantum Unique Drive Control Parameters
3F _H	Return all pages to the initiator (Valid for MODE SENSE command only)

The MODE SENSE data will provide information to the initiator as to how the Q200 Format Parameters and Drive Geometry Parameters are configured if Pages 03_H and 04_H are specified, respectively. The specification of these page codes in a MODE SELECT command will result in a CHECK CONDITION status with ILLEGAL REQUEST sense key.

If Page 3F_H is specified, information from all pages implemented by the Q200 will be returned to the initiator in the MODE SENSE data. This page code is only valid in the MODE SENSE command.

NOTE: If the allocation length is set to 12 bytes or less, the page code field is ignored and illegal page numbers are accepted. If the allocation length is greater than 12 bytes, an illegal page code will be rejected. This allows the initiator to receive the 4-byte header and 8-byte block descriptor without receiving specific page information.

ALLOCATION LENGTH: The allocation length specifies the number of bytes that the initiator has allocated for returned MODE SENSE data. An allocation length of zero indicates that no MODE SENSE data shall be transferred. This condition shall not be considered as an error. Any other value indicates the maximum number of bytes that shall be transferred. The target will terminate the DATA IN phase when allocation length bytes have been transferred or when all available MODE SENSE data have been transferred to the initiator, whichever is less.

MODE SENSE DATA: The MODE SENSE data (Table 6-30) contain a four-byte header, followed by an eight-byte block descriptor, followed by zero or more "pages". The meanings and organization of these data are the same as for the corresponding MODE SELECT data (Section 6.5.10.2), as modified by the option specified in the page control field of CDB byte 2.

6.5.14.2 MODE SENSE DATA

Table 6-30
MODE SENSE DATA

Byte	7	6	5	4	3	2	1	0
0	SENSE DATA LENGTH							
1	MEDIUM TYPE = 0							
2	WP=0	RESERVED = 0						
3	BLOCK DESCRIPTOR LENGTH = 8							

Block Descriptor

0	DENSITY CODE = 00							
1-3	(MSB)	NUMBER OF BLOCKS = 0					(LSB)	
4	RESERVED = 0							
5-7	(MSB)	BLOCK LENGTH					(LSB)	

Page Descriptors

0	PS	R=0	PAGE CODE											
1	PAGE LENGTH													
2 to n	REFER TO PAGE DEFINITION AND SECTION 6.5.10.1													

SENSE DATA LENGTH: The SENSE DATA length specifies the length in bytes of the following MODE SENSE data that is available to be transferred during the DATA IN phase. The SENSE DATA length does not include itself.

WRITE PROTECTED (WP): A write protected bit set to zero indicates that the medium is write enabled. A WP bit set to one indicates that the medium is write protected. For the Q200, the value is always 0.

BLOCK DESCRIPTOR LENGTH: The block descriptor length specifies the length in bytes of all the block descriptors. It is equal to eight for the Q200.

The block descriptor specifies the medium characteristics for the drive. The block descriptor contains a density code, a number of blocks, and a block length, which have the same meanings as corresponding fields in the MODE SELECT parameter list (Section 6.5.10.2).

PARAMETERS SAVEABLE (PS): A parameters saveable bit set to zero in each page header indicates that the supported parameters of that page cannot be saved by the Q200. A PS bit set to one indicates that the supported parameters of the page can be saved by the Q200. All pages with parameters that are changeable by the user are saveable on the Q200.

6.5.14.3 DIRECT ACCESS DEVICE FORMAT PARAMETERS, Page Code 3H:

Byte	Bit	7	6	5	4	3	2	1	0
0		RESERVED = 0							PAGE CODE = 3H
1									PAGE LENGTH = 16H

Handling of Defects Fields

2	(MSB)	TRACKS PER ZONE (0)	
3		TRACKS PER ZONE (4H OR 6H)	(LSB)
4	(MSB)	ALTERNATE SECTORS PER ZONE (0)	
5		ALTERNATE SECTORS PER ZONE (2H)	(LSB)
6	(MSB)	ALTERNATE TRACKS PER ZONE (0)	
7		ALTERNATE TRACKS PER ZONE (0)	(LSB)
8	(MSB)	ALTERNATE TRACKS PER VOLUME (0)	
9		ALTERNATE TRACKS PER VOLUME (0)	(LSB)

Track Format Field

10	(MSB)	SECTORS PER TRACK (0)	
11		SECTORS PER TRACK (20H)	(LSB)

Sector Format Fields

12	(MSB)	DATA BYTES PER PHYSICAL SECTOR (02H)	
13		DATA BYTES PER PHYSICAL SECTOR (0)	(LSB)
14	(MSB)	INTERLEAVE (0)	
15		INTERLEAVE (1H)	(LSB)
16	(MSB)	TRACK SKEW FACTOR (0)	
17		TRACK SKEW FACTOR (AH)	(LSB)
18	(MSB)	CYLINDER SKEW FACTOR (0)	
19		CYLINDER SKEW FACTOR (12H)	(LSB)

Drive Type Field

20	SSEC=1	HSEC=0	RMB=0	SURF=0	INS=0	RESERVED=0
21-23	RESERVED = 0					

The following MODE SENSE pages, 03_H and 04_H, are read only pages available only with the MODE SENSE command. The specification of these page codes in a MODE SELECT command will result in a CHECK CONDITION status with an ILLEGAL REQUEST sense key.

HANDLING OF DEFECTS FIELDS:

The **Tracks Per Zone** field actually indicates the number of tracks per cylinder on the Q200. The value of this field is four for the Q250 and six for the Q280.

The **Alternate Sectors Per Zone** field actually indicates the number of sectors per cylinder that the Q200 shall deallocate from the initiator addressable blocks during the FORMAT UNIT command. These sectors will be available as replaceable sectors for defective sectors. The value of this field is two for the Q200.

The **Alternate Tracks Per Zone** field indicates the number of tracks that the target shall deallocate from the initiator addressable blocks during the FORMAT UNIT command. These tracks will be available as replaceable tracks for defective tracks. The value of zero indicates that the Q200 does not allocate alternate tracks.

The **Alternate Tracks Per Volume** field indicates the number of tracks that the target shall deallocate from the initiator addressable blocks during the FORMAT UNIT command. These tracks will be available as replaceable locations (tracks or sectors) for defect handling. The value of zero indicates that the Q200 does not allocate alternate tracks.

TRACK FORMAT FIELD: The **Sectors Per Track** field indicates the number of physical sectors that the Q200 shall allocate per disk track. This implies that the Q200 set an equal data sector size to all sectors accordingly. The Q200 has 32 sectors per track.

SECTOR FORMAT FIELDS: The **Data Bytes Per Physical Sector** field indicates the number of data bytes that the Q200 shall allocate per physical sector. This value may be different than the logical block size in the Block Descriptor section of the MODE SELECT data. This implies that the Q200 set an equal number of sectors per track accordingly. The Q200 has 512 data bytes per physical sector.

The Q200 has an **interleave factor** of one.

The **Track Skew Factor** field indicates the number of physical sectors between the last logical block of one track and the first logical block on the next sequential track of the same cylinder. The value of this field is 10 for the Q200.

The **Cylinder Skew Factor** field indicates the number of physical sectors between the last logical block of one cylinder and the first logical block on the next sequential cylinder. The value of this field is 18 for the Q200.

DRIVE TYPE FIELD:

An **SSEC** (Soft Sector) bit 7 of byte 20 set to the value of one indicates that the target shall use soft sector formatting. This bit is always set to zero for the Q200 because the Q200 uses hard sector formatting.

An **HSEC** (Hard Sector) bit 6 of byte 20 set to the value of one indicates that the Q200 uses hard sector formatting. The HSEC bit and the SSEC bit are mutually exclusive.

An **RMB** (Removable) bit 5 of byte 20 set to the value of one indicates that the logical unit is removable. An RMB bit set to zero indicates that the logical unit is not removable. This bit is set to zero for the Q200.

The **SURF** (Surface) bit 4 of byte 20 indicates how the Q200 shall map the logical block addressing into physical block addressing. A SURF bit set to the value of one indicates that the target shall allocate progressive addresses to

all sectors on a surface prior to allocating sector addresses to the next surface. A SURF bit set to zero indicates that the Q200 shall allocate progressive addresses to all sectors within a cylinder prior to allocating sector addresses to the next cylinder. This bit is set to zero for the Q200.

The INS (Inhibit Save) bit 3 of byte 20 set to the value of one indicates that the target shall inhibit the saving of any parameters during the successful completion of the next FORMAT UNIT command. The INS bit set to zero indicates that the parameters of the Pages defined by the Page Codes 3 and 4 shall be saved by the Q200 as read only pages. This bit is set to zero for the Q200.

6.5.14.4 RIGID DISK DRIVE GEOMETRY PARAMETERS, Page Code 4H:

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0				PAGE CODE = 4H			
1			PAGE LENGTH = 12H					
2	(MSB)			NUMBER OF CYLINDERS = 0				
3				NUMBER OF CYLINDERS = 3H				
4				NUMBER OF CYLINDERS = 42H			(LSB)	
5				NUMBER OF HEADS (3H or 6 H)				
6-8	(MSB)			STARTING CYLINDER - WRITE PRECOMPENSATION = 0			(LSB)	
9	(MSB)			STARTING CYLINDER - REDUCED WRITE CURRENT = 0				
10				STARTING CYLINDER - REDUCED WRITE CURRENT = 02H				
11				STARTING CYLINDER - REDUCED WRITE CURRENT = 4EH			(LSB)	
12-13	(MSB)			DRIVE STEP RATE			(LSB)	
14-16	(MSB)			LANDING ZONE CYLINDER = 0			(LSB)	
17-19				RESERVED = 0				

The Q200 has 823 cylinders.

The Q250 has four heads; the Q280 has six heads.

The write precompensation starting cylinders, reduced write current starting cylinders, and drive step rate fields are handled by the embedded SCSI controller on the Q200.

The Landing Zone Cylinder field indicates two's complement location where the target shall position the disk heads prior to stopping the spindle with the START/STOP UNIT command. This field is only meaningful with drives which do not automatically seek to the Landing Zone before stopping the spindle. The Q200 automatically parks the heads in the Landing Zone, using Airlock*, upon disconnection of power.

6.5.15 START/STOP UNIT COMMAND: OPCODE = 1B(hex)

Table 6-31
START/STOP UNIT COMMAND

Byte	Bit 7	6	5	4	3	2	1	0		
0	OPCODE = 1B _H									
1	LUN= 0			RESERVED = 0			IMMED			
2-3	RESERVED = 0									
4	RESERVED = 0			START						
5	VU=0		RESERVED = 0			F	L			

The START/STOP UNIT command (Table 6-31) requests that the logical unit be enabled or disabled for further operations.

An **immediate (IMMED)** bit set to one indicates that status will be returned as soon as the operation is initiated. An IMMED bit set to zero indicates that status will be returned after the operation is completed.

A **START** bit set to one requests the drive be started. A START bit set to zero requests that the drive be stopped. A request to start or stop the Q200 may be repeated without error; there will be no effect on the drive if the requested condition is the same as its current condition.

A shorting plug (labeled WS) on the Q200 PCB disables/enables the auto-start option. If the shorting plug is not installed, no command is required to start the drive. If the shorting plug is installed, a START/STOP UNIT command is required to enable the drive for any operations; this configuration is recommended only for systems requiring power sequencing of devices (usually to limit power supply surge currents).

When the Q200 drive is stopped, either due to a "hard" RESET or a START/STOP UNIT command, both the spindle motor and actuator are disabled and do not draw appreciable currents. Starting the drive requires up to 30 seconds before commands requiring disk access (see TEST UNIT READY, Section 6.5.1.1 above) can be executed.

6.5.16 SEND DIAGNOSTIC COMMAND: OPCODE = 1D(hex)

6.5.16.1 SEND DIAGNOSTIC COMMAND STRUCTURE:

Table 6-32
SEND DIAGNOSTIC COMMAND

Bit Byte	7	6	5	4	3	2	1	0		
0	OPCODE = 1DH									
1	LUN= 0			RESERVED = 0		SELF TEST	DEVOFL	UNITOFL		
2	RESERVED = 0									
3-4	(MSB)	PARAMETER LIST LENGTH					(LSB)			
5	VU=0	RESERVED = 0			F		L			

The SEND DIAGNOSTIC command (Table 6-32) requests the drive to perform diagnostic tests on itself. This command is usually followed by a RECEIVE DIAGNOSTIC RESULTS command, except when the self test (SELF TEST) bit is set to one.

The **parameter list length** specifies the length in bytes of the parameter list that shall be transferred during the DATA OUT phase. A parameter list length of zero indicates that no data shall be transferred. This condition shall not be considered as an error.

A **logical unit off-line (UNITOFL)** bit set to one enables write operations on user medium or operations that affect user visible medium positioning. This bit is not supported by Q200 drives. If this bit is set to one on the Q200, the command shall be terminated with a CHECK CONDITION status with the sense key set to ILLEGAL REQUEST with sense code Invalid field.

An **SCSI device off-line (DEVOFL)** bit set to one enables diagnostic operations that may adversely affect operations to other logical units on the same target. This bit is not supported by the Q200. If this bit is set to one on the Q200, the command shall be terminated with a CHECK CONDITION status with the sense key set to ILLEGAL REQUEST with sense code Invalid field.

A **SELF TEST** bit set to one directs the target to complete its default self test. If the self test is requested, the parameter list length shall be set to zero and no data shall be transferred. If the self test successfully passes, the command shall be terminated with a GOOD status; otherwise, the command shall be terminated with a CHECK CONDITION status with the sense key set to HARDWARE ERROR.

6.3.17.2 DIAGNOSTIC PARAMETER LIST:

Vendor Unique, format to be defined.

6.5.17 READ CAPACITY COMMAND: OPCODE = 25(hex)

Table 6-33
READ CAPACITY COMMAND

Byte	7	6	5	4	3	2	1	0		
0	OPCODE = 25H									
1	LUN=0			RESERVED = 0			RELADR=0			
2-5	(MSB) LOGICAL BLOCK ADDRESS					(LSB)				
6-7	RESERVED = 0									
8	VU=0		RESERVED = 0				PMI			
9	VU=0		RESERVED = 0			F	L			

The READ CAPACITY command (Table 6-33) provides a means for the initiator to request information regarding the capacity of the disk drive.

A **partial medium indicator (PMI)** bit set to zero indicates that the information returned in the READ CAPACITY data shall be the logical block address and block length (in bytes) of the last logical block of the logical unit. The logical block address in the command descriptor block shall be set to zero for this option.

A PMI bit set to one indicates that the information returned shall be the logical block address and block length (in bytes) of the last logical block address after which a substantial delay in data transfer will be encountered. This logical block address shall be greater than or equal to the logical block address specified in the command descriptor block.

The eight bytes of READ CAPACITY data shown in Table 6-34 shall be sent during the DATA IN phase of the command.

Table 6-34
READ CAPACITY DATA

Byte	DESCRIPTION		
0-3	(MSB)	LOGICAL BLOCK ADDRESS	(LSB)
4-7	(MSB)	BLOCK LENGTH	(LSB)

Due to the physical organization of sectors and logical block assignments on the Q200, the system designer can predict, to a high level of confidence, the initial LBA for all cylinders. Formatting assigns consecutive logical block addresses to sectors on a cylinder, with rare exceptions. The designer can predict that records stored beginning at a cylinder boundary will probably occupy contiguous physical sectors (skipping occasional defective sectors according to the defect map) and thus be optimally accessible. However exceptions do occur due to grown defects (usually reallocating defective sectors to spare sectors on the same cylinder). Note that when reformatting, grown defective sectors are deallocated along with the factory found defective sectors, using an In-Line sparing method which is described in Appendix A. Therefore the designer is recommended to rely on the READ CAPACITY command to check for optimal file assignment.

To calculate the initial LBA for each cylinder, figure that each sector contains 512 bytes and there are 190 sectors per cylinder for the Q280 (126 for the Q250), and account for 2048-byte logical blocks straddling cylinder boundaries in some cases. Thus initial logical block addresses on the Q280 are:

- 0, 190, 380, 570, . . . for 512-byte blocks;
- 0, 95, 190, 285, . . . for 1024-byte blocks;
- 0, 48, 95, 143, . . . for 2048-byte blocks.

6.5.18 READ EXTENDED COMMAND: OPCODE = 28(hex)

Table 6-35
READ EXTENDED COMMAND

Byte	7	6	5	4	3	2	1	0					
0	OPCODE = 28 _H												
1	LUN = 0			RESERVED = 0			RELADR=0						
2-5	(MSB) LOGICAL BLOCK ADDRESS (LSB)												
6	RESERVED = 0												
7-8	(MSB) TRANSFER LENGTH (LSB)												
9	VU=0	RESERVED = 0			F		L						

The READ EXTENDED command (Table 6-35) requests that the target transfer data to the initiator.

LOGICAL BLOCK ADDRESS: The LOGICAL BLOCK ADDRESS specifies the logical block at which the READ operation shall begin.

TRANSFER LENGTH: The TRANSFER LENGTH specifies the number of contiguous logical blocks of data that shall be transferred. A transfer length of zero indicates that no logical blocks shall be transferred. This condition shall not be considered as an error. Any other value indicates the number of logical blocks that shall be transferred.

6.5.19 WRITE EXTENDED COMMAND: OPCODE = 2A(hex)

Table 6-36
WRITE EXTENDED COMMAND

Byte	7	6	5	4	3	2	1	0		
0	OPCODE = 2AH									
1	LUN = 0			RESERVED = 0			RELADR=0			
2-5	(MSB) LOGICAL BLOCK ADDRESS					(LSB)				
6	RESERVED = 0									
7-8	(MSB) TRANSFER LENGTH					(LSB)				
9	VU=0		RESERVED = 0			F		L		

The WRITE EXTENDED command (Table 6-36) requests that the target write the data transferred by the initiator to the medium.

LOGICAL BLOCK ADDRESS: The LOGICAL BLOCK ADDRESS specifies the logical block at which the WRITE operation shall begin.

TRANSFER LENGTH: The TRANSFER LENGTH specifies the number of contiguous logical blocks of data that shall be transferred. A transfer length of zero indicates that no logical blocks shall be transferred. This condition shall not be considered as an error and no data shall be written. Any other value indicates the number of logical blocks that shall be transferred.

6.5.20 SEEK EXTENDED COMMAND: OPCODE = 2B(hex)

Table 6-37
SEEK EXTENDED COMMAND

Byte	7	6	5	4	3	2	1	0		
0	OPCODE = 2BH									
1	LUN = 0			RESERVED = 0			RELADR=0			
2-5	(MSB) LOGICAL BLOCK ADDRESS					(LSB)				
6-8	RESERVED = 0									
9	VU=0		RESERVED = 0			F		L		

The SEEK EXTENDED command (Table 6-37) requests that the logical unit seek to the specified logical block address.

6.5.21 VERIFY COMMAND: OPCODE = 2F(hex)

Table 6-38
VERIFY COMMAND

Bit Byte	7	6	5	4	3	2	1	0			
0	OPCODE = 2F _H										
1	LUN = 0			RESERVED = 0			BYTCHK	RELADR=0			
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)			
6	RESERVED = 0										
7-8	(MSB)	VERIFICATION LENGTH						(LSB)			
9	VU=0	RESERVED = 0			F		L				

The VERIFY command (Table 6-38) requests that the target verify the data written on the medium.

A byte check (BYTCHK) bit set to zero causes the verification to be simply a medium verification (ECC). This is the only option supported by the Q200. A BYTCHK bit set to one which causes a byte-by-byte compare of data on the medium and the data transferred from the initiator is not supported by the Q200. If this bit is set to one, the command will be terminated with a CHECK CONDITION status with the ILLEGAL REQUEST sense key.

The logical block address specifies the logical block at which the verify operation shall begin.

The verification length specifies the number of continuous logical blocks of data that shall be verified. A transfer length of zero indicates that no logical blocks shall be verified. This condition shall not be considered as an error. Any other value indicates the number of logical blocks that shall be verified.

6.5.22 READ DEFECT DATA COMMAND: OPCODE = 37(hex)

Table 6-39
READ DEFECT DATA COMMAND

Bit Byte	7	6	5	4	3	2	1	0
0	OPCODE = 37 _H							
1	LUN = 0			RESERVED = 0				
2	RESERVED = 0			P	G	DEFECT LIST FORMAT		
3-6	RESERVED = 0							
7-8	(MSB)	ALLOCATION LENGTH						(LSB)
9	VU=0	RESERVED = 0			F		L	

The READ DEFECT DATA command (Table 6-39) requests that the target transfer the medium defect data to the initiator.

The P bit set to one indicates that the initiator requests that the primary (factory) list of defects be returned. The G bit set to one indicates that the initiator requests that the grown list of defects be returned. With bits P and G both set to one, the Q200 is requested to return the primary and the grown list of defects. Only the Defect List Header bytes will be returned if both the P and G bits are set to zero.

The **defect list format** specifies the preferred format for the defect list to be returned by the Q200. If a format not supported by the Q200 is requested, the Q200 will return the defect data in physical sector format. In this case, the Q200 will create the CHECK CONDITION status with RECOVERED ERROR sense key after the defect list data list has been transferred. Two defect data formats are supported by the Q200.

PHYSICAL SECTOR FROM INDEX FORMAT (101_B): This format will return the physical cylinder, head, and sector location from index of the medium defects.

BYTES OFFSET FROM INDEX (100_B): This format will return the physical cylinder, head and the bytes offset from index of the defect location. The offset number reported represents the first byte of the defective sector.

The **allocation length** specifies the number of bytes that the initiator has allocated for returned defect data. An allocation length of zero indicates that no defect data shall be transferred. Any other value indicates the maximum number of bytes that shall be transferred (not counting the defect list header). The target shall terminate the DATA IN phase when allocation length bytes have been transferred or when all available defect data have been transferred to the initiator, whichever is less.

The defect data list contains a four byte header (Table 6-40) followed by zero or more defect descriptors.

Table 6-40
DEFECT LIST HEADER

Byte	7	6	5	4	3	2	1	0
0	RESERVED = 0							
1	RESERVED = 0			P	G	DEFECT LIST FORMAT		
2-3	(MSB)			DEFECT LIST LENGTH			(LSB)	

The bits P and G indicate which defect list is actually returned by the target. The defect list format specifies the actual format of the defect list returned by the Q200. The formats of the defect descriptors are shown in Tables 6-41 and 6-42. The defect list length specifies the total length in bytes of the defect descriptors that follow. The defect list length is equal to eight times the number of defect descriptors. If the allocation length of the CDB is too small to transfer all of the defect descriptors, the defect list length shall not be adjusted to reflect the truncation.

Table 6-41
PHYSICAL SECTOR FORMAT

Byte	Defect Descriptor(s):	
0-2	(MSB)	CYLINDER NUMBER (LSB)
3		HEAD NUMBER
4-7	(MSB)	SECTOR NUMBER (LSB)

Table 6-42
BYTES OFFSET FROM INDEX FORMAT

Byte	Defect Descriptor(s):	
0-2	(MSB)	CYLINDER NUMBER (LSB)
3		HEAD NUMBER
4-7	(MSB)	BYTES OFFSET (LSB)

6.5.23 WRITE BUFFER COMMAND: OPCODE = 3B(hex)

Table 6-43
WRITE BUFFER COMMAND

Byte	Bit 7	6	5	4	3	2	1	0					
0	OPCODE = 3B _H												
1	LOGICAL UNIT NUMBER			RESERVED = 0		MODE							
2	BUFFER ID = 0												
3-5	RESERVED = 0												
6-8	(MSB)	TRANSFER LENGTH				(LSB)							
9	CONTROL BYTE												

The WRITE BUFFER command (Table 6-43) is used in conjunction with the READ BUFFER command as a diagnostic function for testing the target's data buffer memory and the SCSI bus integrity. There is no access to the medium during the execution of this command.

The **buffer control valid (BCV)**, bit 0 of byte 1, must be set to the value of zero for the Q200. This indicates that bytes 2 through 6 are reserved.

The **byte transfer length** specifies the number of bytes to be transferred to and retained in the Q200 data buffer. The byte transfer length contains a four-byte header, followed by the WRITE BUFFER data. A transfer length of zero indicates that no data transfer shall take place. This condition shall not create the CHECK CONDITION status. If the transfer length is greater than the maximum size of the Q200 data buffer, the Q200 shall create the

CHECK CONDITION status with the ILLEGAL REQUEST sense key and with the information bytes reporting the residual count of bytes not transferred. In this case no data shall be transferred from the initiator. It shall not be considered an error to request a transfer length less than the target data buffer size.

6.5.24 READ BUFFER COMMAND: OPCODE = 3C(hex)

Table 6-44
READ BUFFER COMMAND

Bit Byte	7	6	5	4	3	2	1	0					
0	OPCODE = 3CH												
1	LOGICAL UNIT NUMBER			RESERVED		MODE							
2	BUFFER ID = 0												
3-5	RESERVED = 0												
6-8	(MSB)	ALLOCATION LENGTH					(LSB)						
9	CONTROL BYTE												

The READ BUFFER command (Table 6-44) may be used in conjunction with the WRITE BUFFER command as a diagnostic function for testing the target's buffer memory and the SCSI bus integrity. There is no access to the medium during the execution of this command.

The **buffer control valid (BCV)**, bit 0 of byte 1, must be set to the value of zero for the Q200. This indicates that bytes 2 through 6 are reserved.

The **allocation length** specifies the number of bytes that the initiator has allocated for returned data including the header. An allocation length of zero indicates that no data shall be transferred. This condition shall not create the CHECK CONDITION status. Any other value indicates the maximum number of bytes that shall be transferred. This data is to be used by the initiator for comparison with the data pattern sent during the WRITE BUFFER command.

It shall not be considered an error to request an allocation length less than the Q200 data buffer size. The target shall terminate the DATA IN phase when allocation length bytes have been transferred or when all available READ BUFFER data have been transferred to the initiator, whichever is less.

The DATA IN phase of the READ BUFFER command consists of a four-byte header (Table 6-45) followed by the READ BUFFER data.

Table 6-45
READ BUFFER HEADER

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1-3	(MSB)	AVAILABLE LENGTH					(LSB)	

In order to avoid corruption of the target's data buffer by another initiator, it is recommended that the initiator issue the RESERVE UNIT command prior to the WRITE BUFFER, and issue the RELEASE UNIT command after the READ BUFFER command. Any other intervening command, other than REQUEST SENSE, may corrupt the data buffer.

The available length of data bytes returned by the Q200 will be the Q200 maximum buffer size. If the allocation length of the CDB is too small to transfer all of the available length, the available length shall not be adjusted to reflect the truncation.

APPENDIX A MEDIA DEFECT HANDLING

To ensure long-term data integrity, Quantum's Q200 Series Intelligent Disk Drives incorporate media defect handling and error correcting code capabilities. These integrated features eliminate the need for user defect maps and for writing dedicated utility software for defect mapping, thus reducing testing time required during system integration at an OEM site. During system integration and during regular use, the Q200 can continue to scan and compensate for any new defective sectors which may show up later on the disks.

DEFECT DETECTION

Any discrepancy between recorded data and recovered data is defined as a data error. In high density digital recording systems, it is necessary to provide an error detection and correction scheme to enhance the operational performance and increase the reliability of the system.

The errors can be either "soft" -- those that are not readily repeatable, or "hard" -- those that are repeatable with high probability. Soft errors are generally related to the signal-to-noise ratio of the system. They represent marginal conditions of the media, heads, read/write circuitry and the controller data synchronizer circuits. Hard errors are most often due to defects, pits, scratches or thin spots in the media. These defective media areas can be detected and skipped (not used for storage).

DATA ERROR RECOVERY

The integrity of the data stored on the Q200 Series Intelligent Disk Drives is guaranteed by the use of a 96-bit Reed-Solomon Error Correcting Code (ECC). When data is written on the drive, a 12-byte (96 bit) syndrome is generated, appended to the data, and written on the disk along with the data. This syndrome is generated by special hardware on the drive. When data is retrieved from the disk it is again run through the ECC special hardware. If the data is unchanged, the resulting syndrome should be zero. A non-zero syndrome indicates the data has been changed.

The basic strategy of the Q200 during a READ operation is to keep retrying the READ operation until a stable syndrome is found. A stable syndrome means that the same group of data is read twice in a row and the same syndrome is received both times. This strategy prevents invoking correction on soft errors and dramatically reduces the probability of miscorrection. When a stable syndrome is found, the drive will attempt to use ECC to correct the data, assuming ECC correction is enabled via the MODE SELECT Error Recovery Parameters. If the number of retries reaches the maximum specified by the user, the drive will again attempt to correct the error, even if two stable syndromes have not been found (again assuming that ECC correction is enabled).

DEFECT MANAGEMENT

FACTORY FOUND DEFECTS

Q200 Series disk drives are scanned in the factory for defective sectors on the disks. Those defective sectors are deallocated prior to shipment. An In-Line Defect Map, duplicated on tracks on the drive allocated for this purpose, will contain the locations of all defective sectors found in the drive.

Two sectors (for the Q250 and Q280) are allocated per cylinder as replacement sectors. Defects found at the factory are skipped using an "In-Line" sparing method, which replaces each defective sector with the next good sector on the cylinder (Figure A-1). In the rare instance where more than two defects are found within the same cylinder, the additional defects will be mapped into the nearest adjacent cylinder with spare sectors available.

This approach minimizes media used for replacement and reduces access time to the replacement sector. In addition, cylinder boundaries are well defined, with a simple relationship between each initial sector's physical and logical block addresses -- thus providing the host a method to ensure that file records can be stored within cylinder boundaries for optimal accessibility.

Head 0	x	x+1	x+2			x+29	x+30	x+31
Head 1	x+32	(BAD) xxxx	x+33			x+60	x+61	x+62
Head 5	x+159	x+160	x+161				x+188 (spare)	x+189 (spare)

(Q280)
now being
used)

NOTES: x = logical block address
 Q280: cylinder number x 190
 Q250: cylinder number x 126
 (assuming 512-byte logical block size)
 xxxx = deallocated sector
 Zone size is six tracks/zone (default)

Figure A-1 Factory Found Defects

FIELD FOUND DEFECTS

Q200 Series disk drives also contain a Replacement Sector Map in the event that additional defects are found during drive usage. Field found, or "grown" defects are also mapped into the two replacement sectors located at the end of each cylinder. However, field defects are mapped directly into the spare sector locations (without shifting subsequent logical blocks, as in the case of the In-Line sparing method) until the next formatting (Figure A-2). Similar to the factory sparing method, in the rare instance where there are more than two defects per cylinder, the additional defects will be mapped into the nearest adjacent cylinder with spare sectors available. This approach ensures that, in most cases, a long seek to an alternate cylinder will not be necessary in order to manage additional defects found during drive usage.

Head 0	x	x+1	x+2			x+29	x+30	x+31
Head 1	x+32	(BAD) xxxx	x+34			x+61	x+62	x+63
Head 5	x+160	x+161	x+162				x+189 (spare)	x+33 (spare)

(Q280)

NOTES:

- x = logical block address
- Q280: cylinder number x 190
- Q250: cylinder number x 126
- (assuming 512-byte logical block size)
- xxxx = deallocated sector
- Zone size is six tracks/zone (default)

Figure A-2 Field Found Defects

Two software configuration "switches" (Post Error and Automatic Read Reallocation) are available for host control of the field replacement operation. Post Error determines whether or not the Q200 will notify the host every time recovered data is passed from the drive. Automatic Read Reallocation determines whether or not, when corrected data is passed to the host, the Q200 will automatically add the physical location of that sector containing the corrected data onto the Replacement Sector Map, move that sector of data into a replacement sector, then check that the sector was properly duplicated. These two software switches allow the selection of four possible options by the host. The option desired by the host will be specified via the MODE SELECT command.

Option 1 (Post Error off, Automatic Read Reallocation off):

This is a "no action" option. That is, the host will not be informed that recovered data is being passed, nor will the Q200 take any action to map out the suspect sector.

Option 2 (Post Error off, Automatic Read Reallocation on):

This is a transparent replacement option. When recovered data is passed from the Q200, it will not notify the host of the action. However, the drive will automatically add the physical location of the sector containing the corrected data onto the Replacement Sector Map, move that sector of data into a replacement sector, then verify that the move was completed correctly.

Option 3 (Post Error on, Automatic Read Reallocation off):

Every time recovered data is passed to the host, the Q200 will notify the host of the recovered error. The Q200 will not take action to map out the suspect sector.

By logging suspect block addresses passed from the Q200, the host can use the REASSIGN BLOCKS command or the FORMAT UNIT command to manage new defects. The REASSIGN BLOCKS command can be used to add new defective block locations to the Replacement Sector Map and move the data from those sectors into replacement sectors. Alternately, the drive can be reformatted using the logged suspect block addresses with the FORMAT UNIT command (see Reformatting).

Option 4 (Post Error on, Automatic Read Reallocation on):

Every time recovered data is passed to the host, the Q200 will notify the host of the recovered error. The Q200 will automatically add the physical location of that sector containing the corrected data onto the Replacement Sector Map, move that sector of data into a replacement sector, and verify that the move was completed correctly.

Reformatting

The FORMAT UNIT command issued from the host will cause the Q200 to reorganize data on the disk in a manner identical to the way it was formatted in the factory (i.e., using an In-Line sparing method). In addition to using the original factory In-Line Defect Map and the Field Replacement Sector Map, the reformatting operation can also be performed using a host-supplied list of defective logical blocks.

APPENDIX B DISCACHE

To enhance drive performance, Quantum's Q200 Series Disk Drives incorporate DisCache, a 60K-byte look-ahead disk cache. This integrated feature is user programmable, using the MODE SELECT command, and can significantly improve system throughput. Since the cache parameters are programmable, the cache can be dynamically configured in real time to best meet overall system requirements.

DISCACHE DESCRIPTION

DisCache uses 60 KB of RAM to anticipate host-system requests for data and to store that data for faster access. When the host requests a particular segment of data, the caching feature uses a prefetch strategy to "look-ahead" and automatically store the subsequent data from the disk into RAM. If the host requests this subsequent data, then the high-speed RAM is accessed rather than the disk.

Since typically 50 percent or more of all disk requests are sequential, there is a high probability that subsequent data requested will be in the cache. This cached data can be retrieved in microseconds rather than milliseconds. As a result, DisCache can provide substantial time savings during half or more of all disk requests. In these instances, DisCache may save most of the disk transaction time by eliminating the seek and rotational latency delays that dominate the typical disk transaction. For example, in a 1 KB data transfer, these delays comprise 80 percent of the elapsed time.

DisCache works by continuing to fill its cache memory with adjacent data after transferring data requested by the host. Unlike a non-caching controller, Quantum's disk controller continues a read operation after the requested data has been transferred to the data buffer. This read operation terminates after a programmed amount of subsequent data has been read into the cache memory. When the 60 KB cache memory becomes full, then the cache contents are replaced using a first-in, first-out strategy.

During write operations, DisCache employs a write-through design, where data written to the drive are written immediately. This design ensures data integrity in case of a power loss. Data written to the disk drive are always stored immediately on the non-volatile disk surface.

The cache memory is 60K bytes of a 64K byte DRAM buffer allocated to hold the data directly accessed by the host via the READ and WRITE commands. The memory functions as a ring buffer with a rollover point at 60K. The unit of data stored is the logical block (i.e., a multiple of the 512 byte sector). Therefore, all accesses to the cache memory must be in multiples of the sector size. In those cases where the cache memory must be used for scratch memory, as in the case of a FORMAT command, or where the size of the logical block may change, as in the MODE SELECT command, the cache will be emptied. The commands that will force emptying of the cache are:

COPY	FORMAT UNIT
INQUIRY	MODE SELECT
MODE SENSE	READ DEFECT LIST
	REASSIGN BLOCKS

PERFORMANCE BENEFITS

In a drive without DisCache, during sequential reads, there would be a delay due to rotational latency even if the disk actuator were already positioned at the desired cylinder. DisCache eliminates this rotational latency time -- 8 milliseconds on average -- when requested data resides in the cache.

Moreover, the disk often must service requests from multiple processes in a multitasking or multiuser environment. In these instances, while each process may request data sequentially, the disk drive must share time among all these processes. In most disk drives, the heads must move from one location to another. With

DisCache, even if another process interrupts, the Quantum drive continues to access the data sequentially from its high-speed memory. In handling multiple processes, DisCache achieves its most impressive performance gains, saving both seek and latency time -- 34 milliseconds on average -- when desired data resides in the cache.

MAXIMUM FLEXIBILITY, EASY TO USE

Quantum designed this feature to be extremely flexible because cache performance is highly application-dependent. DisCache features programmable options that enable users to adjust caching parameters to optimize performance. These options can be specified and subsequently modified using an interface software command -- SCSI's MODE SELECT command. The following table outlines the programmable parameters and how they can be used.

Table B-1
Programmable Parameters

Programmable Parameters	DisCache Function	Input Range	Default Value
Cache enable/disable	Activates cache when set (bit=1)	0,1	1
Number of Cache Segments	Sets number of cache segments to be maintained	1-12	12
Write Indexing Enable/Disable	Allows data from write operation to be overwritten in the cache memory to increase cache utilization when this bit is disabled (=0).	0,1	1
Maximum Prefetch Absolute	Sets maximum amount of data to be prefetched as an absolute amount or as a multiple of the size of the read request.	0-116 logical blocks	0
Maximum Prefetch Multiple		0-63 logical blocks	3
Minimum Prefetch Absolute		0-116 logical blocks	0
Minimum Prefetch Multiple	Both minimum and maximum amounts can be set.	0-63 logical blocks	0
Prefetch Threshold	Disables prefetch operation if read request exceeds this threshold; this avoids prefetching during large read operations and increases cache utilization.	0,1	1

For more detailed information on these parameters, and how to program them, please see Section 6.5.10.5 "Quantum Unique Cache Control Parameters".

Through the use of these programmable parameters, Quantum allows for the tailoring of its caching feature to optimize individual system performance. At the same time, DisCache can be completely transparent to the system or end user.

DESIGNED FOR HIGH PERFORMANCE

DisCache itself was designed with high performance in mind by incorporating both programmable options and low overhead. DisCache is configured by Quantum to guarantee that system disk requests are handled as a first priority and look-ahead caching occurs only during the drive's idle time. Specifically, the look-ahead prefetch will immediately terminate if the initiator sends a subsequent command before the maximum amount has been prefetched. Efficient cache management ensures minimum command processing time (a few hundred microseconds) associated with this feature.

Finally, the disk drive is the ideal location for this type of look-ahead cache because here the least amount of system overhead is involved. System-level caches incur the same amount of overhead in reading requested data as in anticipating requests. In either case, the system must access the disk and delay other disk accesses as well as other operations. With DisCache, anticipating requested data occurs during otherwise idle time and does not delay other disk accesses. From the system perspective, the overhead cost is minimal while the benefit may be quite significant.

